END
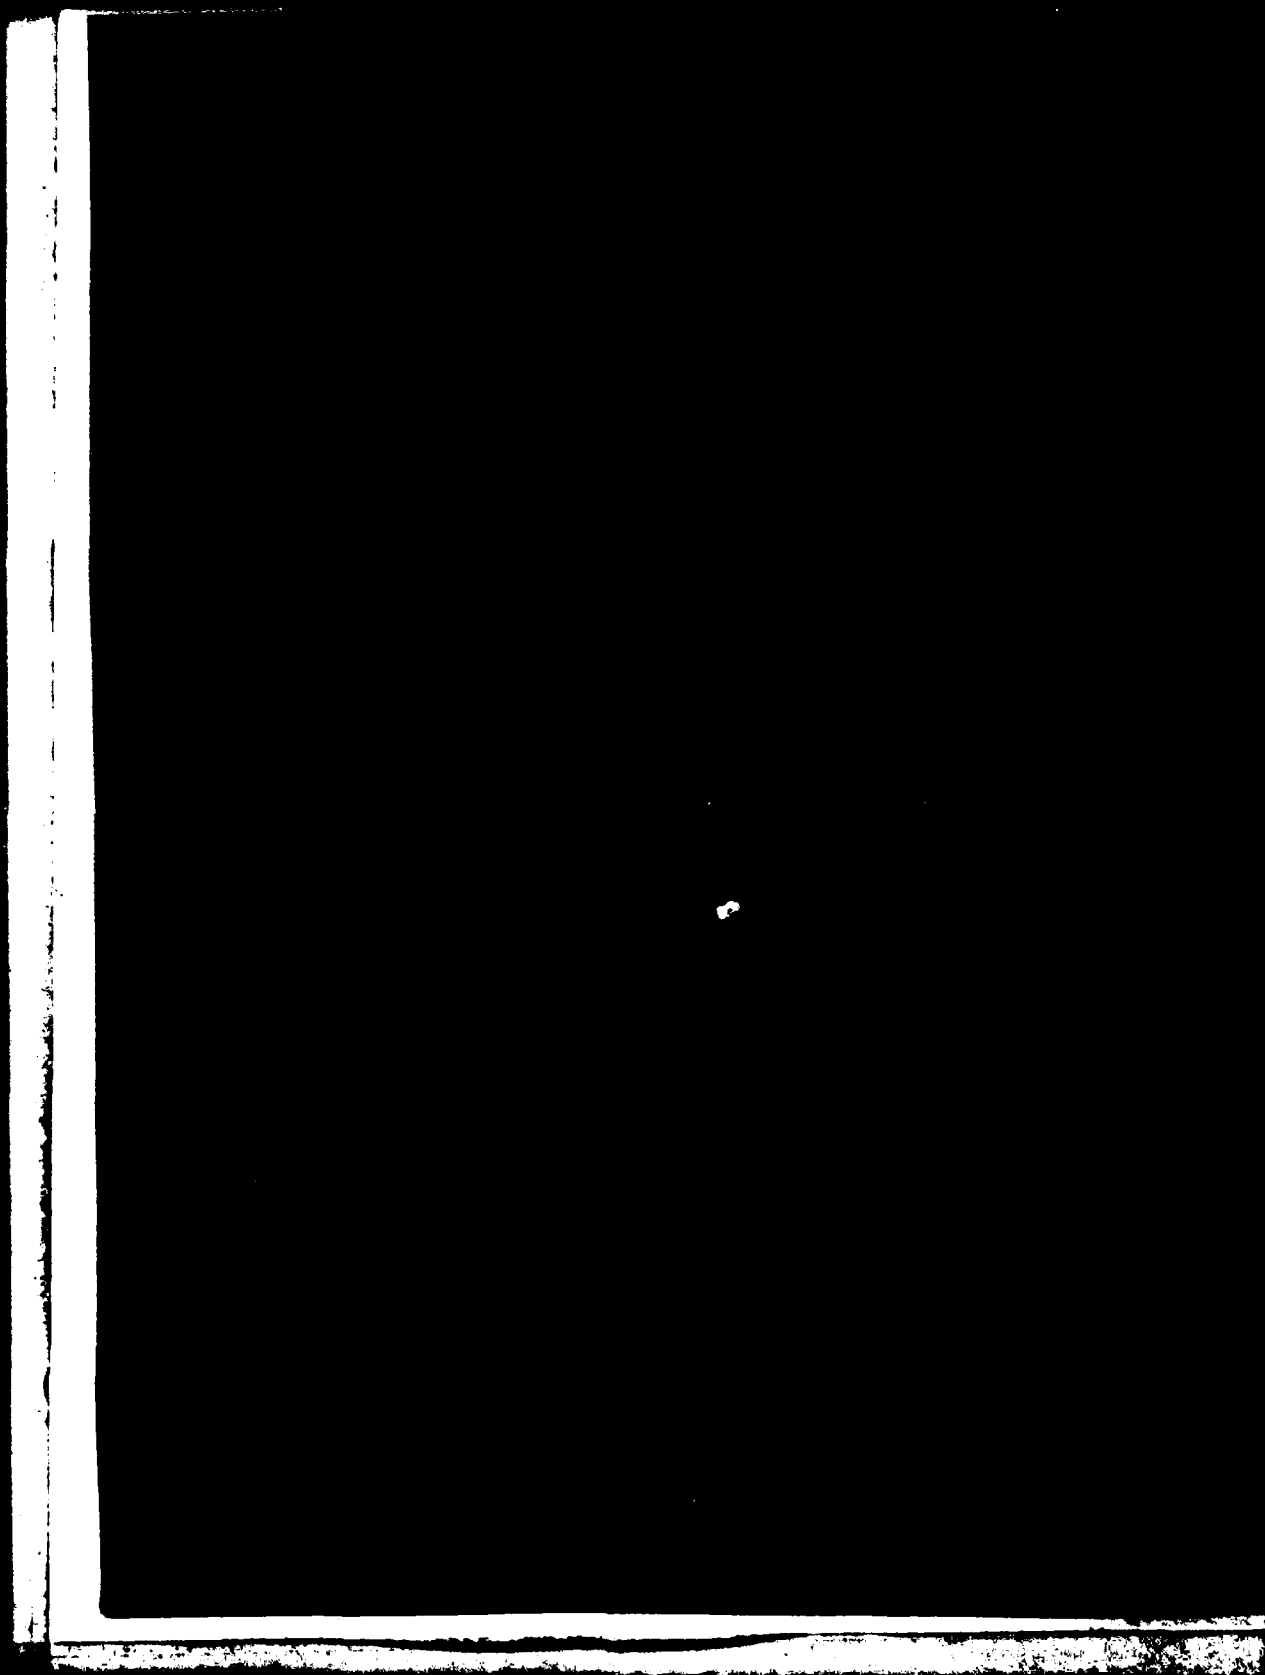DATE
FILMED

DTIC

MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS-1963-A

MASSACHUSETTS INSTITUTE OF TECHNOLOGY
LINCOLN LABORATORY

DEFENSE SWITCHED NETWORK TECHNOLOGY
AND EXPERIMENTS PROGRAM

ANNUAL REPORT
TO THE
DEFENSE COMMUNICATIONS AGENCY

1 OCTOBER 1981 — 30 SEPTEMBER 1982

ISSUED 5 APRIL 1983

DTIC
ELECTE
S JUN 7 1983
D
A

LEXINGTON   MASSACHUSETTS

ABSTRACT

This report documents work performed during FY 1982 on the DCA-sponsored Defense Switched Network Technology and Experiments Program. The areas of work reported are: (1) development of routing algorithms for application in the Defense Switched Network (DSN); (2) instrumentation and integration of the Experimental Integrated Switched Network (EISN) test facility; and (3) EISN experiment planning, coordination, and execution.

iii

TABLE OF CONTENTS

LIST OF ILLUSTRATIONS

LIST OF TABLES

DEFENSE SWITCHED NETWORK
TECHNOLOGY AND EXPERIMENTS PROGRAM

1. INTRODUCTION AND SUMMARY

This report documents work performed during FY 1982 on the DCA-sponsored Defense Switched Network Technology and Experiments Program. The areas of work reported are: (a) development of routing algorithms for application in the Defense Switched Network (DSN); (b) instrumentation and integration of the Experimental Integrated Switched Network (EISN) test facility; and (c) EISN experiment planning, coordination, and execution.

Routing algorithm development efforts during FY 82, described in Sec. 2, have focused on: (a) further development of the new mixed-media routing and preemption procedures, which were originally introduced in FY 81; (b) evaluation of steady-state routing algorithm performance using a network analysis program; and (c) development of a new call-by-call network simulation program to evaluate dynamic routing algorithm performance including the effect of Multi-Level Precedence and Preemption (MLPP) features. Extensive performance evaluation of the new routing algorithms showed that they substantially improved network performance after damage. For example, the most effective new algorithm, referred to as adaptive-mixed-media routing, provided an improvement in service to the most poorly served users slightly greater than the improvement associated with adding 10-percent more land trunks in a 20-node test network. The call-by-call simulator includes all new routing procedures, MLPP features, and common-channel signaling (CCS) user-level protocols compatible with the CCITT No. 7 standard. The simulator has been tested extensively, and will be utilized for routing performance evaluation tests during FY 83.

During FY 82, Lincoln Laboratory has continued its major role in the development and integration of experimental subsystems for EISN. This effort is described in Sec. 3. Experimental subsystems developed by the Laboratory include: (a) a Packet/Circuit Interface (PCI) allowing access to the satellite channel from a circuit switch on a T1 carrier format digital interface; (b) a Telephone Office Emulator (TOE) which performs digital switch functions necessary for testing with the PCI; and (c) an Internet Packet Gateway (IPG) allowing connection between packet data networks and the WB SATNET for data protocol experiments. Major FY 82 accomplishments in EISN instrumentation and integration have included:

1

(a) integration, deployment, and test of PCIs and TOEs at the Defense Communications Engineering Center (DCEC) and Lincoln;

(b) construction of PCI and TOE hardware for the three additional EISN sites;

(c) development and test of a Terrestrial Alternate Routing (TAR) capability;

(d) integration and test of IPGs with the Exploratory Data Network (EDN) at DCEC and with the ARPANET at Lincoln.

FY 82 efforts in EISN experiment planning, coordination, and execution are described in Sec. 4. An EISN Experiment Plan and an FY 82 Work Plan detailing long- and short-term experiment planning were prepared and delivered to DCEC in FY 82. Lincoln has taken on an expanded role in EISN system coordination, including initial implementation of a plan for the Laboratory to become the frequency and power level reference station for the network. EISN experiments during FY 82 focused on two-node (DCEC and Lincoln) experiments in satellite/terrestrial integration, routing, and internet data communication. Experiment configurations and milestones are described in Sec. 4.3. Finally, detailed design and planning efforts have begun on an advanced routing/control experimental facility including commercial switches and outboard routing/control processors (RCPs).

2

2. ROUTING ALGORITHM DEVELOPMENT FOR THE DSN

2.1 INTRODUCTION AND SUMMARY

The Defense Communications Agency (DCA) has developed a plan for a new circuit-switched network to eventually replace current systems and become the primary voice and data telecommunications network for the Department of Defense. This new network, called the Defense Switched Network (DSN), must meet the dual objectives of survivable communications under stressed conditions and economical service under normal operations. To support these objectives, the DSN will utilize a mix of transmission media, including both broadcast satellite and point-to-point terrestrial links. Routing procedures are needed to effectively utilize this mix of media. Lincoln efforts during FY 81 and FY 82 have included a major focus on the development of routing procedures for application in the DSN, to meet the dual objectives of survivability and economy.

During FY 81 (see Ref. 1), three new classes of routing procedures were developed for application in the DSN, referred to as (a) mixed-media routing, (b) adaptive-mixed-media routing, and (c) precedence flooding. During FY 82, research has focused on (a) further development of new routing and preemption procedures, (b) evaluation of steady-state routing algorithm performance using a modified network analysis program, and (c) development of a call-by-call network simulation program to evaluate dynamic routing algorithm performance including the effects of Multi-Level Precedence and Preemption (MLPP) features. The first two topics are covered in some detail in a Technical Report[2] which has been delivered separately to the DCA. Here, we shall review the major results presented in Ref. 2, omitting many details and covering some additional items, particularly in the area of preemption.

The report given here on routing algorithm development for the DSN is organized as follows. The next few paragraphs briefly summarize efforts and results in (a) new routing and preemption procedures, (b) steady-state performance evaluation, and (c) call-by-call simulator development. Section 2.2 is a descriptive overview of the new routing and preemption procedures which have been developed. Sections 2.3 and 2.4 describe steady-state performance evaluation techniques and results. Finally, Sec. 2.5 describes the development of the call-by-call simulator.

2.1.1 New Routing and Preemption Procedures

All routing procedures developed for the DSN treat satellite and terrestrial links separately and uniquely, both when routing tables are created and when calls are routed. All procedures also use common-channel signaling to pass call-setup information between switches. This information includes a list of all switches and satellites already in the call path, a list of unavailable earth stations and satellites, and satellite hop and link limits. *Mixed-media* routing procedures use fixed routing tables and three different call processing rules (spill forward control, remote earth station querying, and single-stage crankback). *Adaptive-mixed-media* routing procedures adapt routing tables when parts of the network are destroyed. *Precedence flooding* procedures route high-precedence calls using flooding techniques, and low-precedence calls using mixed-media procedures. A new two-stage precedence flooding routing procedure was developed during the past year. High-priority calls are first routed using one of the mixed-media procedures. If a call is blocked, it is then rerouted using flooding techniques. Otherwise, only mixed-media routing is used. This procedure greatly reduces the required common-channel signaling bandwidth; however, it still finds a path to the destination if the path exists. A new preemption procedure, called guided preemption, was also developed over the past year. This procedure finds the shortest preemptable path to the destination and then preempts the fewest lower-precedence calls on that path. It is designed to preempt fewer low-priority calls than the blind preemption technique used in AUTOVON, or source-destination preemption as proposed in Ref. 3.

2.1.2 Steady-State Performance Evaluation

The steady-state performance of (a) mixed-media routing with spill-forward control, (b) mixed-media routing with remote earth station querying, and (c) adaptive-mixed-media routing was determined using a modified Defense Communications Engineering Center (DCEC) steady-state network analysis program. Details of the steady-state analyses are available in a Lincoln Laboratory Technical Report[2] that has been delivered to DCEC. The new routing procedures were compared with modified forward routing and primary-path-only routing. Evaluations were performed using 20- and 40-node mixed-media networks under overload, with various patterns of offered traffic, and with different amounts and types of network damage. MLPP features were not included in the comparison because of the difficulty of incorporating these features in the trunk group queueing theory model used in the analysis program.

4

The new routing procedures studied, especially adaptive-mixed-media routing, substantially enhanced network performance after damage. These procedures did not reduce the average point-to-point blocking probability. They did, however, improve the service provided to the most poorly served group of users, and they denied the possibility of call completion to the fewest users. The improvement in service to the most poorly served users provided by adaptive-mixed-media routing was slightly greater than the improvement associated with adding 10-percent more land trunks. Under overload conditions and when offered traffic patterns were dramatically shifted, the new procedures performed as well as or better than the best of the other procedures which were studied.

2.1.3 Call-by-Call Simulator Development

A new call-by-call network simulator is being developed to evaluate the performance of routing procedures when MLPP features are available. This simulator will include all routing procedures including precedence flooding and mixed-media routing with crankback. It will also include three types of preemption: (a) blind, as in AUTOVON; (b) source-destination, as proposed by General Telephone and Electronics (GTE) in Ref. 3; and (c) guided, a new type of preemption developed as part of this year's program. The simulator includes common-channel signaling (CCS) user-level protocols compatible with CCITT No. 7. It monitors the number of bits sent over the CCS network and also the call-setup-time delay caused by transmitting CCS messages over limited-rate CCS links. It also produces output printouts and plots of: point-to-point blocking for five precedence levels, number of calls preempted, CCS bits transmitted, call-setup time, number of links per call, number of remote earth station query CCS messages, number of crankback CCS messages, link occupancy, and other statistics on traffic flow and network performance. The input and output of the simulator are compatible with those of current DCEC steady-state analysis programs, and the simulator can be run either at Lincoln Laboratory or DCEC. The simulator has been tested with all routing procedures except precedence flooding, and can be run both with and without blind preemption. We are currently continuing testing and debugging, and plan to add the two new types of preemption and precedence flooding.

2.2 OVERVIEW OF ROUTING AND PREEMPTION PROCEDURES FOR MIXED-MEDIA NETWORKS

2.2.1 Mixed-Media Networks

All new routing and preemption procedures are designed for mixed-media networks which have key characteristics in common with proposed characteristics of the DSN.

Mixed-media networks are circuit-switched networks with many small programmable switches that exchange routing and control information over a common-channel signaling network. Mixed-media networks also include both point-to-point and broadcast connectivity that could be obtained using terrestrial trunks and broadcast satellite connection as in the DSN. Alternatively, it could be obtained using high-frequency radio broadcast transmission and microwave and high-frequency radio point-to-point links. Figure 1 illustrates a small mixed-media network with one point-to-point satellite link connecting node 1 to node 4, and one broadcast DAMA satellite connecting nodes 9, 15, and 16. Implicit in this figure is a common-channel signaling packet-switched network which, in the simplest case, parallels the network of voice trunks.
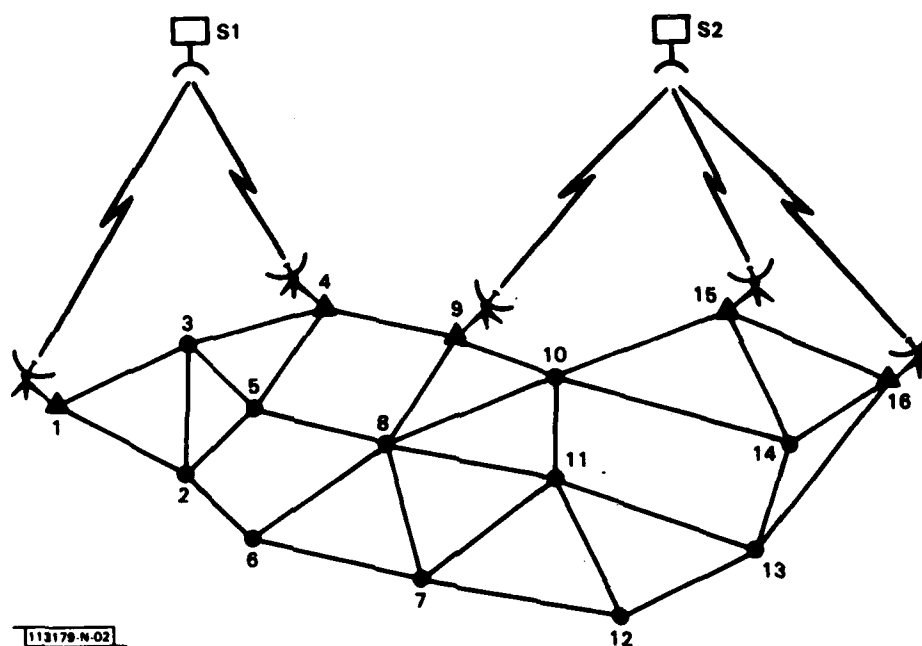
Fig. 1. A mixed-media network with 2 satellites, 5 earth stations, and 16 nodes.

2.2.2 Routing and Preemption Procedures

The objectives of new routing and preemption procedures developed for use in mixed-media networks are: (a) to provide economical performance under normal operation, (b) to provide desired performance after various types of network damage, and (c) to minimize the number of low-priority calls preempted. Details of these new procedures are available in

6

Refs. 1 and 2. Brief descriptions of three new types of routing procedures (mixed-media routing, adaptive-mixed-media routing, and precedence flooding) follow. A new preemption procedure, called guided preemption, is also described.

A. Mixed-Media Routing

The simplest routing procedures which have been developed for use in mixed-media networks are called mixed-media procedures. These procedures use fixed routing tables and call-processing rules which employ either spill-forward control, a new type of control called remote earth station querying, or crankback. A major distinguishing characteristic of mixed-media procedures is that satellite and land links are treated separately and uniquely, both when routing tables are created and when calls are actually routed through a network. For example, call-processing rules use information about the status of key nodes with satellite earth stations to route calls. In addition, routing tables indicate whether the shortest path to the destination switch via a given outgoing link includes a satellite hop. The use of fixed routing tables and local signaling in these procedures enhances routing security and automatically protects unaffected parts of a network when intentional or unintentional signaling errors occur in specific locations. These procedures also tend to minimize signaling and switch CPU processing bandwidth requirements, and to minimize switch and signaling hardware requirements in general.

Call-processing rules used with mixed-media routing require a common-channel signaling network in which switches automatically sense failure or destruction of attached links, and adjacent switches, earth stations, and satellites. Call-request messages are sent over this network to establish call paths. These messages contain a special header that includes a trace or list of switches and satellites currently in the call path. The header also includes a list of known blocked or destroyed earth stations and satellites, the maximum number of links allowed in the land and satellite routing to the destination, and the maximum number of satellite hops allowed. Limits vary for different precedence-level calls and for voice and data calls. Information in the call-request header is used by switches to prevent loops and shuttles (a loop between two nodes), to route calls away from or avoid blocked or destroyed earth stations and satellites, to limit the length of alternate routes, to prevent excessive delay, and to direct the cranking back of calls.

Spill-Forward Control: — The simplest call-processing rule for mixed-media routing is called spill-forward control which either blocks a call at a switch or routes a call to another

7

SATELLITE AVAILABLE

SATELLITE BLOCKED,
SPILL FORWARD

●●● SATELLITE BLOCKED,
REMOTE EARTH STATION
QUERYING

S1

123432-N-02

ROUTING TABLE IN NODE 4
FOR DESTINATION NODE 14

LINK	USE SATELLITE?	EARTH STATION
4-3	YES (S1)	3
4-5	YES (S1)	3
4-2	YES (S1)	3
4-8	NO	–

ROUTING TABLE IN NODE 3
FOR DESTINATION NODE 14

LINK	USE SATELLITE?	EARTH STATION
3-11	YES (S1)	3
3-4	NO	–
3-5	NO	–

ROUTING TABLE IN NODE 1
FOR DESTINATION NODE 14

LINK	USE SATELLITE?	EARTH STATION
1-2	YES (S1)	3
1-4	YES (S1)	3
1-7	NO	–
1-4	NO	–
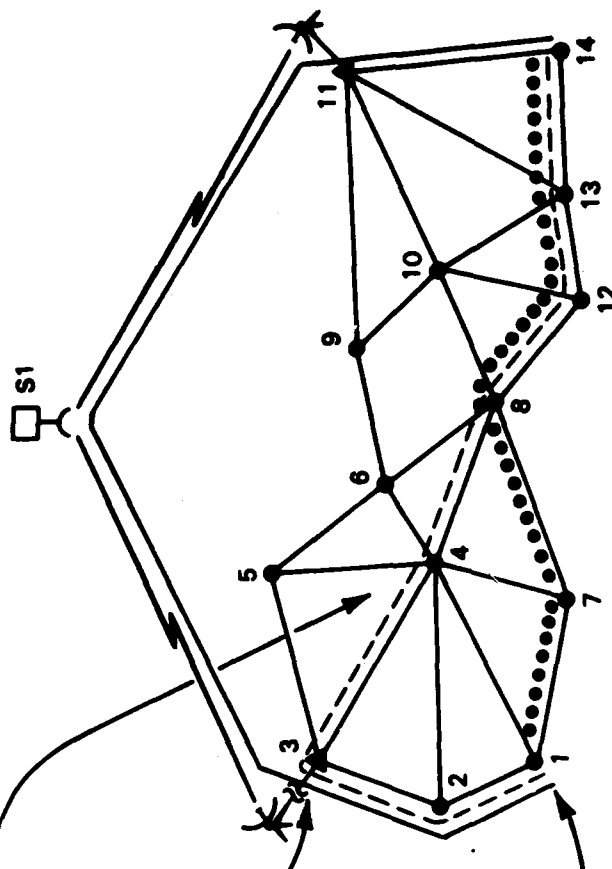
Fig. 2. Call paths through a 14-node mixed-media network using mixed-media routing and either spill-forward control or remote earth station querying.

8

switch not yet in the call path. The next switch in the call path is selected by sequentially examining the links listed in the routing table. A link is skipped if (1) it leads to a switch already in the call path, (2) its associated shortest path includes any of the blocked or destroyed earth stations and satellites listed in the call-request header, (3) no more satellite hops are allowed and the link's associated shortest path includes a satellite, (4) no free or preemptable trunks are available on the link, or (5) the link is destroyed. A call is blocked if no acceptable routing-table entry is found after examining as many entries as is allowed for the call's precedence level, or if only one more link is allowed in the call path and the next switch is not the destination.

Call paths established using spill-forward control are illustrated in Fig. 2. The mixed-media network in this figure includes 14 switches, 2 earth stations, and 1 satellite. The routing tables shown for destination node 14 include ordered lists of outgoing links and also information on the first earth station and satellite on the shortest path to node 14 via each link.

The solid line in Fig. 2 illustrates the call path from node 1 to node 14, established when all links are free and the satellite is free. Starting with switch 1, each switch places itself on the trace in the call-setup message header and routes the call using the first routing-table entry. When the call request arrives at node 14, the trace includes 1-2-3-S1-11 and the list of unavailable earth stations and satellites is empty.

The dashed line in Fig. 2 illustrates the call path when the satellite is busy. The call-request message travels to switch 3 where the first routing-table entry is blocked. Switch 3 recognizes this, places the earth station in node 3 onto the list of unavailable earth stations in the call-request header, and routes the call to node 4 using the second routing-table entry, 3-4. The switch at node 4 skips the first three routing-table entries because they lead to shortest paths which include the blocked earth station at node 3 (this earth station is on the list in the call-request header). The switch at node 4 routes the call to node 8 using the last routing-table entry, 4-8. All other nodes route the call using the first entry in their routing table for destination node 14. When the call request arrives at node 14, the trace includes 1-2-3-4-8-12-13, and the list of unavailable earth stations includes the earth station at node 3. Note that, if the list of unavailable earth stations had not been passed to switch 4, the call would have been routed from switch 4 to switch 5 on link 4-5, the first routing-table entry which does not lead to a node in the trace. The call would then have been routed to nodes 6 and 8 and then blocked and lost due to too many links in the call path.

9

ROUTING TABLE IN NODE 2
FOR DESTINATION NODE 14

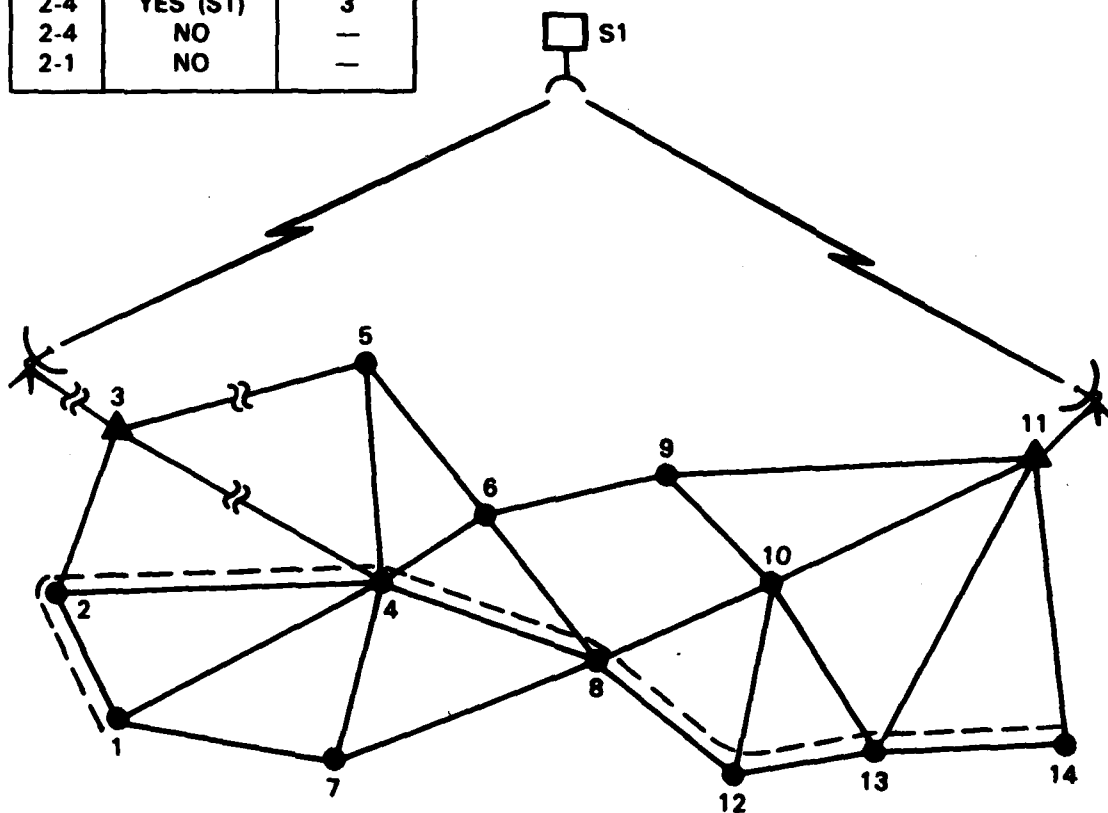| LINK	USE SATELLITE?	EARTH STATION
2-3	YES (S1)	3
2-4	YES (S1)	3
2-4	NO	—
2-1	NO	—

Fig. 3.   Call path from node 1 to node 14 when mixed-media routing with single-stage crankback is used, earth station at node 3 is blocked, and links between nodes 3 and 4 and between nodes 3 and 5 are blocked.

Remote Earth Station Querying: — The detour in the dashed curve of Fig. 2 can be avoided by a new type of call processing called remote earth station querying which is made possible by common-channel signaling (CCS). Whenever the shortest path to a destination includes a satellite, a CCS query message is sent to the first earth station on this path to determine the status of the earth station and associated satellite. If the earth station and satellite are available, the call is routed normally as with spill-forward control. If not, the earth station or satellite is added to the list of unavailable earth stations and satellites stored in the call-request header, and then the call is routed normally as with spill-forward control.

The dotted curve in Fig. 2 illustrates the call path with remote earth station querying when the earth station at node 3 is blocked. The detour associated with spill-forward control has been completely eliminated. With remote earth station querying, switch 1 (the first switch in the call path) initially sends a query message to the earth station at node 3 because this earth station is on the shortest path to the destination (see first routing-table entry). The return message from switch 3 indicates that the earth station is blocked. Switch 1 thus adds the earth station to the list of unavailable earth stations in the call-request header and routes the call on link 1-7, the first routing-table entry not associated with a path that includes the earth station at node 3. The other switches also route calls using routing-table entries not associated with paths that include the earth station at node 3. When the call request arrives at node 14, the trace includes 1-7-8-12-13 and the list of blocked earth stations includes the earth station at node 3.

Single-Stage Crankback: — Single-stage crankback is the third type of call processing used with mixed-media routing. This type of processing is similar to spill-forward control except that a call request blocked at a switch is routed backwards to the previously visited switch which then tries other untried outgoing links. This results in many more alternate paths than are available with spill-forward control. Paths are limited, however, to prevent loops and shuttles, limit path lengths, and prevent routing to blocked earth stations. These limits are needed to prevent excessively long alternate routes which could degrade network performance under heavy traffic loads.

Figure 3 illustrates the call path when the earth station at node 3 is blocked, links 3-4 and 3-5 are blocked, and single-stage crankback is used. The call-request message travels first to node 3. Here, no outgoing links are available and the call would be blocked if spill-forward control were used. Single-stage crankback, however, allows the call to crank back to node 2. Before it is cranked back, switch 3 adds the earth station at node 3 to the list of unavailable

11

earth stations in the *call-request header*. When the call request arrives back at node 2, it thus indicates that the earth station at node 3 is busy. This prevents the call from being routed toward node 3 and causes the call to be routed over land links to node 14. The trace when the call arrives at node 14 includes 1-2-3-2-4-8-12-13, and the list of unavailable earth stations includes the earth station at node 3. Note that the call path would have been identical to that illustrated by the dashed link in Fig. 2 if link 3-4 had been free.

B. Adaptive-Mixed-Media Routing

Adaptive-mixed-media routing procedures are identical to the previously described static-mixed-media procedures under normal operating conditions. When the network is damaged, however, routing tables are automatically updated to enhance survivability.

The procedures used to update routing tables in adaptive-mixed-media routing are similar to those currently used in the ARPANET.[4] Each node stores global information describing network topology. This information is updated only when the network topology changes (switches or links are added, removed, or destroyed). Updated information is transmitted from nodes which detect a change using flooding on a secure CCS network. Routing tables are recomputed in each node when an update is received via the algorithm used to generate the original set of routing tables. This procedure is relatively simple, but it provides adaptive, distributed routing based on global information, and it does not suffer from any inherent adaptation rate, stability, or convergence problems that could occur in procedures which adapt continuously on the basis of network performance.

The main disadvantage of adaptive-mixed-media routing is the extra complexity in switches that it requires and the necessity of transmitting global information describing topological changes around the network.

C. Precedence Flooding

Precedence flooding routes high-priority traffic using flooding techniques alone or a combination of mixed-media routing backed up by flooding. Low-priority calls are routed using mixed-media routing. The flooding scheme is tailored to mixed-media networks. Whenever the destination of an originating high-priority call is more than one link away, the originating switch sends SEARCH messages out over all attached links with free or preemptable trunks. SEARCH messages are transmitted over the CCS network. Each switch adds information to the SEARCH message and forwards the new message to adjacent switches. The

modified SEARCH message is not sent back over the link used by the incoming SEARCH message, and it is not forwarded over links containing no free or preemptable trunks.

Switches forward the first SEARCH message received and any following SEARCH message which comes over a shorter but not necessarily quicker path. This prevents SEARCH messages delayed on a satellite hop from being discarded when land routes are also available. The information contained in SEARCH messages is updated at each switch and includes the number of satellite hops and links traversed, the number of links where preemption is required, and an overall measure of link loading on the SEARCH message path. Switches store this information and the incoming link for as many as N SEARCH messages received for each call, where N is an empirically selected number. Information is stored only for those N SEARCH messages with the shortest distance to their source.

The path-selection process begins at the destination switch. This switch waits a short pre-specified time after receiving the first SEARCH message for a given call, and then attempts to route the call over that link which leads to the path with the fewest links. If paths including the same number of links are available, then link loading, delay, and the number of preempted calls on each path are taken into consideration. If the link selected no longer includes a free trunk, then the link with the next shortest path to the destination is chosen. This process is initiated at tandem switches by a backwards call-setup message which starts at the destination and stops at the source. Tandem switches prevent loops and shuttles in the call path by skipping over links which lead to nodes already in the call path. These switches also reserve trunks on the selected links.

After the backwards call-setup message reaches the source, a forward call-setup message is sent over the established path and the call begins. A call is blocked if the originating switch times out and does not receive a backwards call-setup message after a given period of time.

Precedence flooding is attractive because the shortest path is actually found by an exhaustive search. Its intuitive appeal is also supported by the good performance of flooding in Ref. 5. Flooding is reasonable in the DSN, however, only for high-priority traffic because of its large signaling and switch CPU processing requirements. Conservative assumptions indicate that 2400-baud, common-channel signaling links (obtained, for example, from single-voice trunks) and PDP-11/34-type switch controllers could support, at most, only the projected high-priority AUTOVON traffic for 1985 (roughly 1300 erlangs). The use of

4800-baud common-channel signaling links would increase the allowable traffic which can be routed using flooding to 2500 erlangs. More than 1300 erlangs could be supported by 2400-baud, CCS links if flooding is only used for high-priority traffic when no path is available with mixed-media routing. This two-stage routing plan is very attractive because flooding is used only when it is needed and CCS traffic is greatly reduced relative to pure flooding.

D. Guided Preemption

A new preemption algorithm, referred to as guided preemption, was developed this past year. When preemption is necessary, guided preemption finds the shortest preemptable path between two switches and then preempts as few low-precedence calls as possible on that path. Call-request messages are first routed to the destination using any of the new routing procedures. If preemption is required on only one link in the call path, then the longest-duration call in progress is preempted on that link when the "call-request-successful" message travels back to the source. If preemption is required on two or more links, then switches in the call path exchange information describing the identity of preemptable low-precedence calls. Those calls with the most links in common with the call path are then selected and preempted.

Guided preemption differs from blind preemption used in AUTOVON in that (1) low-precedence calls are not preempted if the call-request message is blocked before it reaches the destination, and (2) the minimum number of calls is preempted. This is illustrated in Fig. 4. which shows a network with 5 switches and 6 low-priority calls in progress. If a high-precedence call is routed from switch 1 to switch 5 over path 1-2-4-5, it is possible to preempt the three low-precedence calls D, E, and F with blind preemption. Guided preemption preempts only call B. If call B isn't available, guided preemption preempts only calls C and F. Guided preemption also preempts no calls if the call-request message travels from switch 1 to 2 to 4 and then is blocked at switch 4 because there are no free or preemptable trunks on link 4-5. Blind preemption could preempt the two low-priority calls D and E in the same situation without successfully completing the high-precedence call.

Guided preemption differs from source-destination preemption described in Ref. 3 in that (1) if a complete path to the destination is not available, the minimum number of lower-precedence calls are preempted and (2) calls are routed over the shortest preemptable paths to the destination. Consider Fig. 4; as a call is being routed from switch 1 to switch 5. If

14
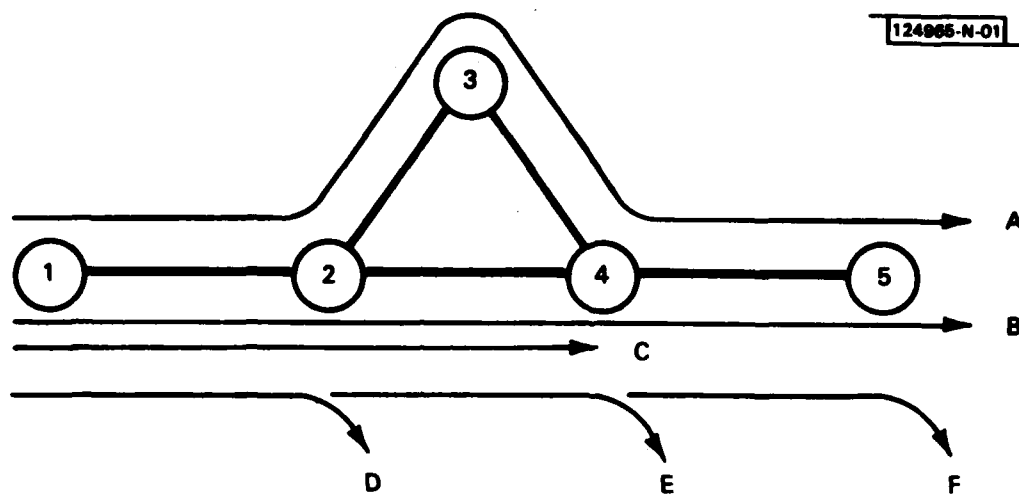
Fig. 4. Network with 5 switches and 6 low-priority calls in progress (A, B, C, D, E, F).

preemption is required on link 1-2, source-destination preemption first tries to preempt calls going through switch 1 that originate or terminate at switch 5. In this case, source-destination preemption could choose call A even though preempting call A and following its path results in a longer call path than would result if call B were chosen. Guided preemption always preempts call B. In addition, if calls A and B are not available to preempt, guided preemption preempts only the two calls C and F. Source-destination preemption may preempt the three calls D, E, and F.

2.3 STEADY-STATE PERFORMANCE EVALUATION TECHNIQUES

The performance of two simple reference routing procedures and three new mixed-media procedures was compared using a steady-state network analysis program and 20- and 40-node test networks with various traffic conditions. Routing procedures, networks, the analysis program, and analysis conditions used in this comparison are described in this section.

2.3.1 Routing Procedures

Two simple reference routing procedures were used — primary-path-only routing and modified forward routing.[6] Primary-path-only routing provides a baseline performance limit

15

obtained without alternate routing when calls can be routed on only one path to each destination. Modified forward routing is a previously used variant of the simplest non-hierarchical routing procedure which allows alternate routes. That procedure, forward routing, routes calls only to switches which are closer to the destination than the current switch. Modified forward routing differs from forward routing in that a call may be routed backwards once to a switch directly connected to the destination via a land link or a satellite.

The three mixed-media routing procedures used were (a) mixed-media routing with spill-forward control, (b) mixed-media routing with remote earth station querying, and (c) adaptive-mixed-media routing. Other new procedures were not included because the performance of these procedures could not be determined using the queueing theory model of trunk group operation incorporated in the analysis program.

2.3.2 Test Networks

Two 20-node and two 40-node, mixed-media networks were created to evaluate routing procedures. Characteristics of these networks are presented in Table I. Link and switch locations of networks DSN1 and DSN2 are presented in Figs. 5 and 6. Solid lines in these figures represent land links, and dashed lines represent links to the one DAMA satellite. Networks DSN1 and DSN2 are minimum-cost networks without any spare trunking to improve survivability, while networks DSN3 and DSN4 include spare terrestrial trunks which cost roughly 10-percent more than the trunks in DSN1 and DSN2. These trunks provide supplementary paths when the satellite is destroyed.

All network designs were based on the projected offered traffic and node locations in a tentative 100-node DSN configuration which was under investigation at the DCEC. The 20 and 40 nodes with the most originating traffic were selected, and earth stations that access one DAMA satellite were positioned on those nodes with the most offered traffic. Land links and link capacities in networks DSN1 and DSN2 were assigned using a minimum-cost network design program[7] modified to allow a DAMA satellite. The cost assigned to a satellite hop in this program was selected iteratively to route roughly one-third of all the offered traffic over the satellite and the desired link-blocking probability was set to 0.1. Networks DSN3 and DSN4 were designed by removing the satellite and earth stations from DSN1 and DSN2 and repeating the design procedure. New long-distance terrestrial links were created

16

TABLE I

CHARACTERISTICS OF FOUR TEST NETWORKS

Network	DSN1	DSN2	DSN3	DSN4
Switches	20	40	20	40
DAMA Satellites	1	1	1	1
Earth Stations	5	9	5	9
Terrestrial Links	118	355	121	361
Terrestrial Voice Trunks	1992	4620	2120	4852
Satellite Capacity (Voice Trunks)	450	930	450	930
Offered Traffic (erlangs)	1450	2960	1450	2960
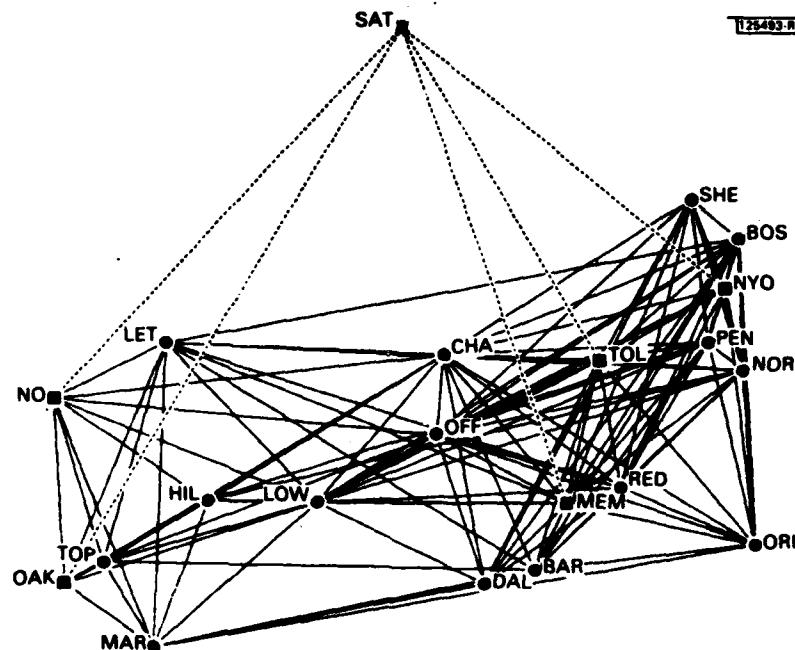
Fig. 5. Test network DSN1.

during this design to compensate for the missing satellite capacity. The largest, most efficient, long-distance links were added to the original networks one at a time until the cost of all terrestrial trunks in the new networks exceeded the cost of the trunks in the original networks by 10 percent.

2.3.3 Analysis Program

A steady-state network analysis algorithm first presented by Katz[8] was used to determine network performance. A program which implemented an updated version of the algorithm developed by Fischer and Knepley[9] and by Fischer et al.[7] was provided by DCEC and modified to include the new routing procedures. The analysis algorithm, described in detail in Ref. 9, is an iterative procedure with a traffic distribution and a link-queueing-theory analysis phase in each iteration. This algorithm was modified by adding mixed-media routing with spill-forward control to the first traffic distribution phase of each iteration, and by modifying the second phase to allow a DAMA satellite. Controls used in the traffic distribution phase were added to block calls with paths that are more than one link longer than the shortest land path to the destination.
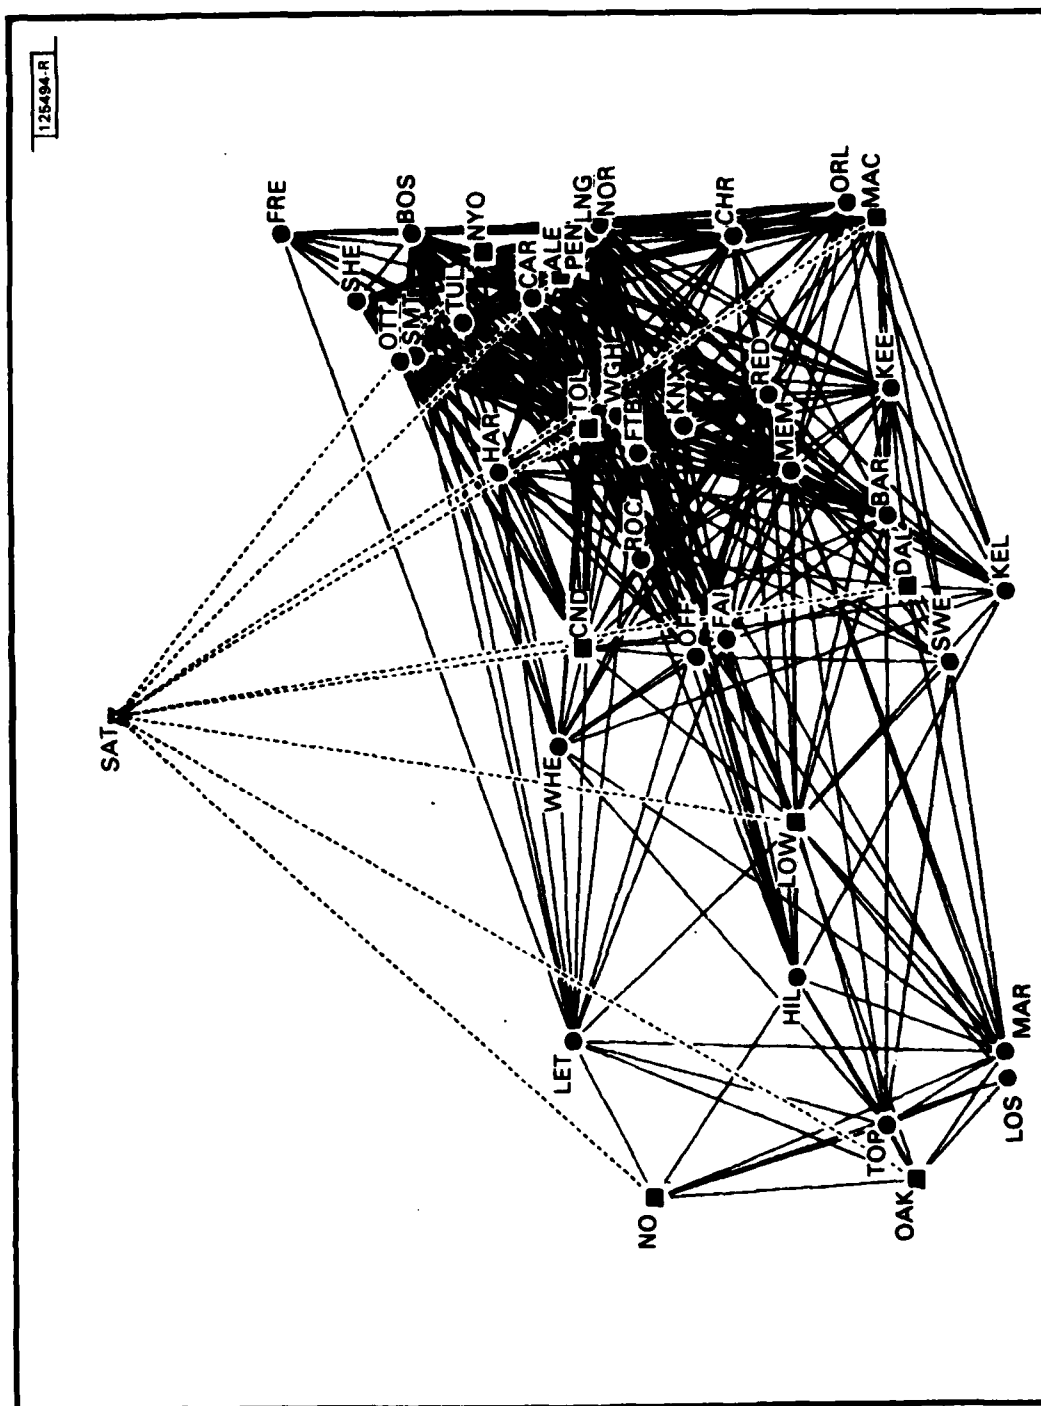
Fig. 6.   Test network DSN2.

19

A new routing-table generation program was written for the modified analysis program. The metric used to order paths in this program is the total number of links biased by total distance in miles. A satellite hop was counted as one link, and the distance assigned to a satellite hop was adjusted to route roughly one-third of all traffic over the satellite. Whenever possible, routing tables included ten entries for each source-destination pair.

2.3.4 Analysis Conditions

Networks were analyzed under normal operating conditions and (a) with different types of network damage, (b) when the traffic was uniformly increased and decreased, and (c) when traffic patterns between source-destination pairs were varied in a number of ways while the total offered traffic was held constant.

Four patterns of network damage were examined. First, from 10 to 40 percent of all terrestrial voice trunks were destroyed. The largest terrestrial links were destroyed one at a time until the desired percentage of trunks was destroyed. Second, the satellite and from 10 to 40 percent of all terrestrial trunks were destroyed. Third, from 2 to 8 of the switches which originate and terminate the most offered traffic were destroyed. Fourth, the satellite and from 2 to 8 switches were destroyed.

Network performance was examined with four deviant traffic patterns. The first was an increase in the offered traffic between all node pairs of 5 to 25 percent. This examines the behavior of routing procedures under uniform overload. The second traffic pattern randomly varied the offered traffic between each node pair from zero to twice the normal value using four different randomizations. The third and fourth traffic patterns varied traffic over large unexpected extremes. The third pattern applied an identical amount of traffic between each node pair but maintained the total offered traffic at a normal value. The fourth pattern (REVERSE MAX/MIN) applied the maximum offered traffic between any node pair to that node pair which normally has the minimum offered traffic, the second largest-offered traffic between any node pair to that node pair with the second least-offered traffic, etc.

2.4 STEADY-STATE PERFORMANCE RESULTS

2.4.1 Performance Under Conditions of Network Damage

The average traffic-weighted point-to-point blocking probability in network DSN1, under normal operation and when the satellite and 0 to 40 percent of all land trunks are destroyed, is presented in Fig. 7. Under normal operation, all routing procedures except
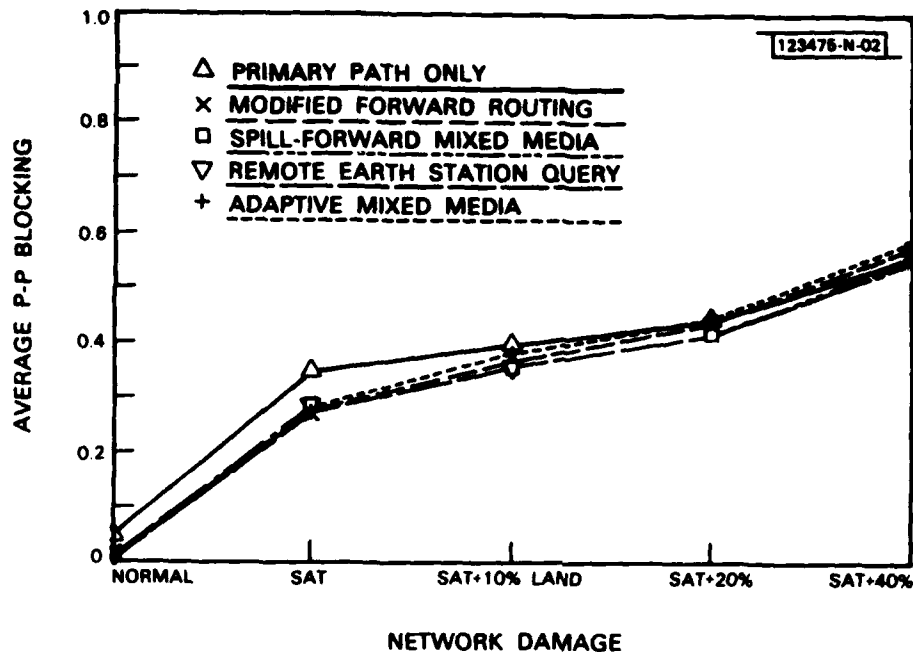
Fig. 7. Average point-to-point blocking probability under normal operation and when the satellite and then 10 to 40 percent of largest land links are destroyed in network DSN1.

primary-path-only routing provide an average blocking of less than 0.02. After the satellite is destroyed, blocking increases substantially. This increase is followed by further increases as more and more land trunks are destroyed. Blocking is generally highest for primary-path-only routing, especially when only the satellite is destroyed and all traffic normally routed over the satellite (roughly one-third of all offered traffic) is blocked. However, there are no major differences between the other routing procedures. Similar results were obtained with all networks. These results are misleading because they are based only on average blocking probabilities. It is important to examine point-to-point blocking probabilities between all nodes in a network.

Histograms of traffic weighted point-to-point blocking probabilities in DSN1 for primary-path-only routing and adaptive-mixed-media routing when the satellite is destroyed are presented in Fig. 8(a). These histograms were created by first placing the offered traffic between every node pair in the bin corresponding to that node pair's point-to-point blocking probability. The total traffic in each bin was then normalized to be a percentage of the total

21

Fig. 8(a). Histograms of traffic weighted point-to-point blocking probabilities between all node pairs when the satellite is destroyed in network DSN1, and either primary-path-only or adaptive-mixed-media routing is used.

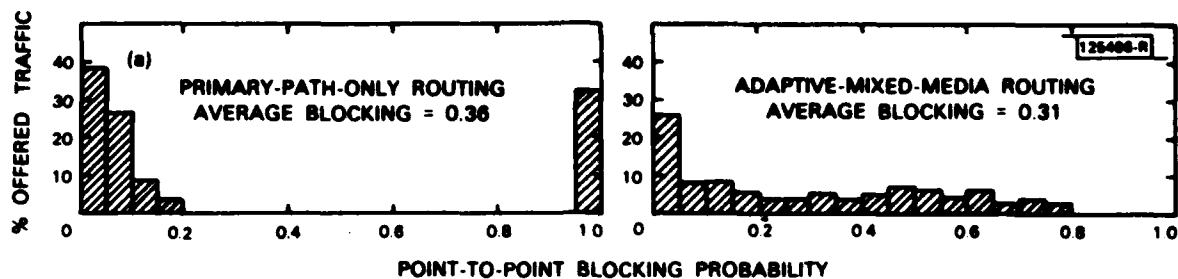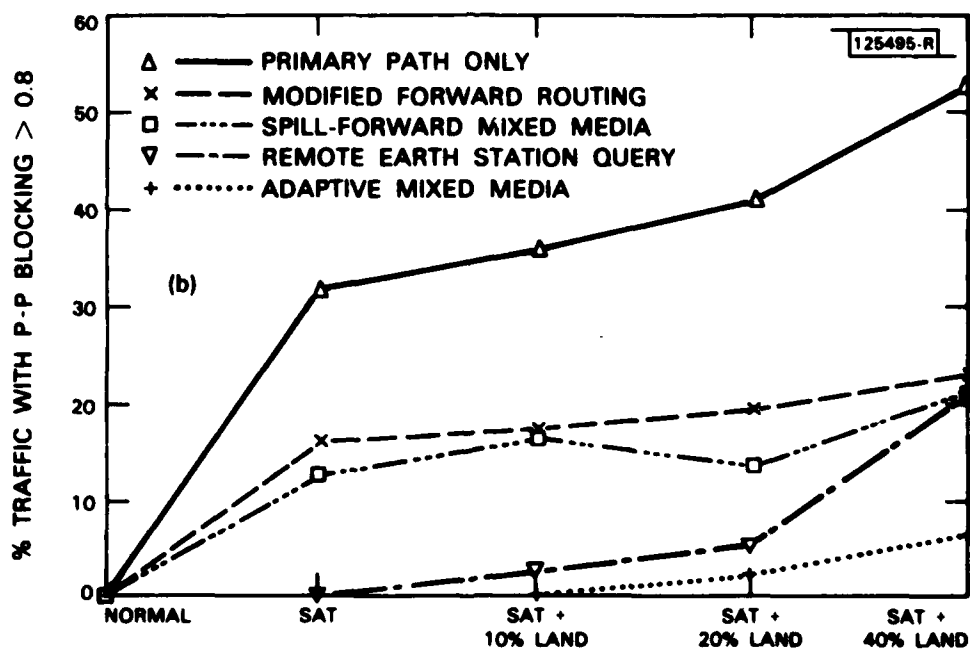Fig. 8(b). Percentage offered traffic which experiences a point-to-point blocking probability >0.8 when the satellite and then 10 to 40 percent of largest land links are destroyed in network DSN1.

22

offered traffic. For example, the histogram for adaptive-mixed-media routing indicates that roughly 25 percent of all offered traffic experiences a blocking of 0.0 to 0.05, that roughly 9 percent of all offered traffic experiences a blocking of 0.05 to 0.1, etc. Under normal conditions, all traffic experiences a blocking probability less than 0.05 and the bin extending from 0.0 to 0.05 contains 100 percent of the traffic.

A comparison of the histograms in Fig. 8(a) indicates that, although the average blocking probabilities of the two routing procedures are not dramatically different, the treatment of the node-to-node pairs with the worst service is extremely different. Roughly 31 percent of all offered calls experience a blocking probability greater than 0.95 (in this case 1.0) with primary-path-only routing. It is impossible to place a call between those node pairs in the bin ranging from 0.95 to 1.0. Users placing calls between those pairs (roughly 31 percent of all users) are treated unfairly and provided unsatisfactory service. Adaptive-mixed-media routing, however, treats all nodes more uniformly and doesn't block communication between any nodes. The highest blocking between any two nodes is less than 0.8 and it is possible, but sometimes difficult, to place a call between all nodes.

The average blocking is similar in Fig. 8(a) because traffic which is completely blocked with primary-path-only routing is eliminated from the network. More resources are left for the remaining traffic which experiences low blocking. Adaptive-mixed-media routing doesn't completely block any traffic. All traffic competes for limited resources, and the range of blocking probabilities experienced is larger. Primary-path-only routing in this case is similar to a congestion control mechanism which blocks all long-distance calls at their source. Although such controls may minimize the average blocking, they are unacceptable because they deny service to critical users by unfairly penalizing users placing long-distance calls.

A comparison between routing procedures based on the percentage of offered traffic provided with unacceptable service is presented in Fig. 8(b). Here it is assumed that service is unacceptable if the probability of blocking on the first attempt of a call is greater than 0.8. The large differences between routing procedures evident in the traffic weighted histograms is clearly evident. When the satellite is destroyed, service is unacceptable for roughly 31 percent of all traffic with primary-path-only routing, but service is acceptable for all traffic with adaptive-mixed-media routing. Service is acceptable for all traffic under normal operation with all routing procedures. The percentage of traffic provided with unacceptable service, however, increases rapidly when the satellite is destroyed and then more slowly as land trunks are destroyed. Primary-path-only routing is clearly unacceptable. From 31 to

23

60 percent of all traffic is provided inadequate service when the network is damaged. Adaptive-mixed-media routing provides the best performance, followed by remote earth station query-ing, mixed-media routing with spill-forward control, and modified forward routing. Similar results were obtained with network DSN2. Differences are greatest when only the satellite is destroyed. Here in DSN1, both remote earth station querying and adaptive-mixed-media routing provide acceptable service for all traffic, service is unacceptable for 13 to 16 percent of all traffic with spill-forward mixed-media routing and modified forward routing, and ser-vice is unacceptable for 31 percent of all traffic with primary-path-only routing. The advan-tage provided by adaptive-mixed-media routing is maintained as land trunks are destroyed. The advantage of remote earth station querying diminishes as more and more land links are destroyed and no new information about this destruction is obtained. In addition, spill-forward mixed-media routing performs slightly better than modified forward routing under almost all conditions.

Improved performance with adaptive-mixed-media routing and remote earth station querying was also obtained when the satellite and then nodes were destroyed in all networks. Adaptive-mixed-media routing also provided best performance when land trunks or nodes were destroyed Under all these conditions, average point-to-point blocking probabilities dif-fered little over routing procedures. The percentage offered traffic with point-to-point block-ing greater than 0.8 was always lowest with adaptive-mixed-media routing, followed by remote earth station querying. Percentages were always highest with primary-path-only rout-ing, and were intermediate with modified forward routing and spill-forward mixed-media routing.

A comparison of the effectiveness of (a) adding 10-percent more trunks to a network to that of (b) using a more complex routing procedure is presented in Fig. 9. The percentage offered traffic with unacceptable blocking is plotted for network DSN1 when the satellite is destroyed. With primary-path-only routing, roughly one-third of all traffic experiences unacceptable blocking. This percentage drops to roughly 16 percent with modified forward routing. After this reduction, the percentage can be reduced by either adding more trunks or changing to adaptive-mixed-media routing. Adding 10-percent more trunks (switching to network DSN3 but still using modified forward routing) reduces the percentage of traffic experiencing unacceptable blocking to 8.5 percent. Changing to adaptive-mixed-media rout-ing, however, reduces the percentage of traffic experiencing unacceptable blocking to zero.

24
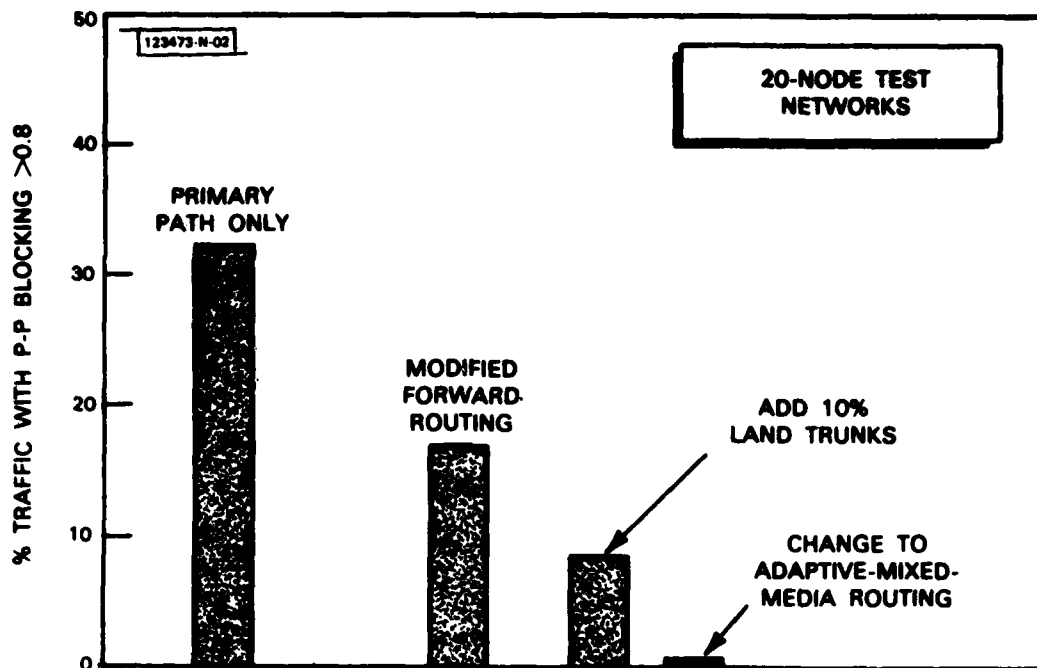
Fig. 9. Percentage offered traffic that experiences a point-to-point block-
ing probability >0.8 when the satellite is destroyed (1) in network DSN1 with
primary-path-only routing, (2) in network DSN1 with modified forward
routing, (3) in network DSN3 which has 10-percent more land trunks than in
DSN1 but still with modified forward routing, and (4) in network DSN1
with adaptive-mixed-media routing.

This result illustrates alternative methods of providing survivability. Spare, expensive trunk-
ing can be provided, or a sophisticated routing procedure can be used. In this example,
increasing the trunking cost by 10 percent was not as effective as implementing adaptive-
mixed-media routing. Similar results were obtained with networks DSN2 and DSN4. The
percentage traffic with unacceptable blocking in DSN2 was 33 percent with primary-path-
only routing and 17 percent with modified forward routing. Adding 10-percent more land
trunks reduced the percentage to 8 percent, and changing to adaptive-mixed-media routing
reduced it to 7 percent.

25

2.4.2 Performance with Uniform Traffic Overload and with Chaotic Traffic

Results obtained with uniform traffic overload in network DSN1 are presented in Fig. 10. At normal loads, the average blocking with modified forward routing and spill-forward mixed-media routing is similar and much lower than the average blocking with primary-path-only routing. At overloads, blocking with modified forward routing and spill-forward mixed-media routing increases rapidly until it equals and slightly exceeds blocking with primary-path-only routing. Similar results were obtained in networks DSN2, DSN3, and DSN4. These results demonstrate that the new routing procedures which provide better performance after network damage do not degrade performance with traffic overload. This is an important result because the increased routing flexibility provided by complex routing procedures could lead to excessively long path lengths under traffic overload which reduce network utilization and degrade performance. It is evident from Fig. 10 that this doesn't happen with mixed-media routing. This result is due, in part, to controls that explicitly limit the number of links in a call path to N links more than the number of links in the shortest path to the destination.
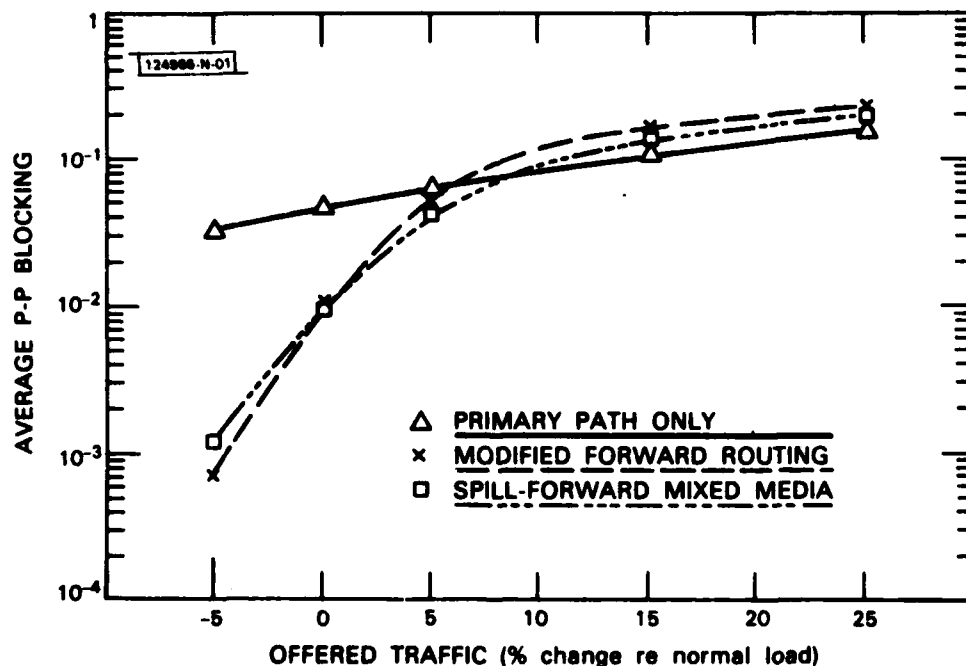

Fig. 10. Average point-to-point blocking with 5-percent underload and 0- to 25-percent uniform traffic overload in network DSN1.

26

With chaotic traffic patterns, there was little difference between modified forward routing and the new routing procedures in either the average point-to-point blocking or the percentage blocking greater than 0.8. Primary-path-only routing, however, performed worse than the other routing procedures.

2.4.3 Summary of Results

A summary of the steady-state analysis results is presented in Table II. A plus in this table indicates that a routing procedure performed better than modified forward routing, and a minus indicates worse performance. A double plus indicates significantly better performance. The performance measures used to generate this table were the average point-to-point blocking probability and the percentage offered traffic with unacceptable blocking (blocking >0.8). As can be seen, adaptive-mixed-media routing provided the best performance under all damage conditions. Remote earth station querying provided good performance when the satellite and then nodes or links were destroyed. All new procedures performed better than modified forward routing when the satellite and then nodes or links were destroyed. All new procedures also performed slightly better than modified forward routing under overload. No routing procedure performed better than modified forward routing with chaotic traffic patterns. All new routing procedures improved performance under damage and overload, and did not degrade performance with chaotic traffic. These new procedures, especially adaptive-mixed-media routing, are thus viable candidates for the DSN.

2.5 PERFORMANCE EVALUATION VIA CALL-BY-CALL SIMULATION

A call-by-call simulator of mixed-media networks using new routing procedures is needed in addition to the steady-state analysis program to (a) examine the effect of MLPP features, (b) examine the dynamic performance of routing procedures, (c) examine the performance of precedence flooding and mixed-media routing with crankback, (d) examine CCS traffic, (e) test CCS user-level protocols, and (f) examine call setup times.

After examining all existing simulators, a decision was made to write a new simulator specifically for this program. Work on the simulator started in the second quarter of FY 82, and it can now be used with all new routing procedures except precedence flooding both with or without blind preemption. The simulator is written in a modern structured language

27

TABLE II

RESULTS OF STEADY-STATE ANALYSIS

Routing Procedure	Test Condition						
	Network Damage					Traffic Overload	Chaotic Traffic
	Satellite	Satellite and Links	Satellite and Nodes	Links	Nodes		
Primary Path Only	−	−	−	−	−		−
Modified Forward Routing	+					+	
Spill-Forward Mixed Media	++	+	+			+	
Remote Earth Station Query	++	++	++			+	
Adaptive Mixed Media	++	++	++	++	++		

called RATFOR which is automatically translated into highly portable FORTRAN code. To date, over 140 pages of RATFOR code have been written, and the simulator has been tested extensively.

A block diagram of the simulator is presented in Fig. 11. Input files include controls for the current run (routing procedure, type of preemption, type and time of damage, duration of run, etc.), network topology, and trunk group sizes. A call generator program uses this information to create a list of offered calls, and a routing table generation program uses the information to create routing tables. Offered calls, routing tables, and the original input files are read into the main simulator DRIVER/MONITOR program. This program reads in calls one at a time and offers them to the network, initiates damage, terminates calls, monitors calls, and prints output reports.

The "Subscriber Loop and CCS" module behaves like a subscriber loop on an originating or terminating switch, and like a CCS network between switches. It transports CCS messages between switches, keeps statistics on CCS traffic and computes CCS transit time, keeps overall statistics, outputs call details, and transports subscriber loop signals (new call, phone ringing, phone answered, phone busy) between originating and terminating switches and telephones.

The "Routing and MLPP Processor" module performs the routing and preemption functions that are necessary in a switch. It is purposely written to be easily modified for placement in a real switch. This processor accepts a CCS message, decides what action to take, and then outputs one or more new CCS messages. It is the most complex part of the simulator and includes both call processing and routing and preemption logic. It implements all new routing and preemption procedures and also new user-level CCS protocols compatible with CCITT No. 7. It uses routing tables, information on the status of all trunk groups, and information on calls in progress to perform these functions.

Outputs of the simulator include: (a) a description of the network and controls; (b) periodic reports on blocking probability, path length, call setup time, reasons for blocking, CCS bits transmitted, preemption statistics, crankback statistics, and remote earth station querying statistics; (c) damage statistics; (d) overall reports on link statistics including blocking and occupancy; (e) overall reports on point-to-point blocking probability between all node pairs; (f) node statistics including traffic and blocking to and from each node; (g) path-length
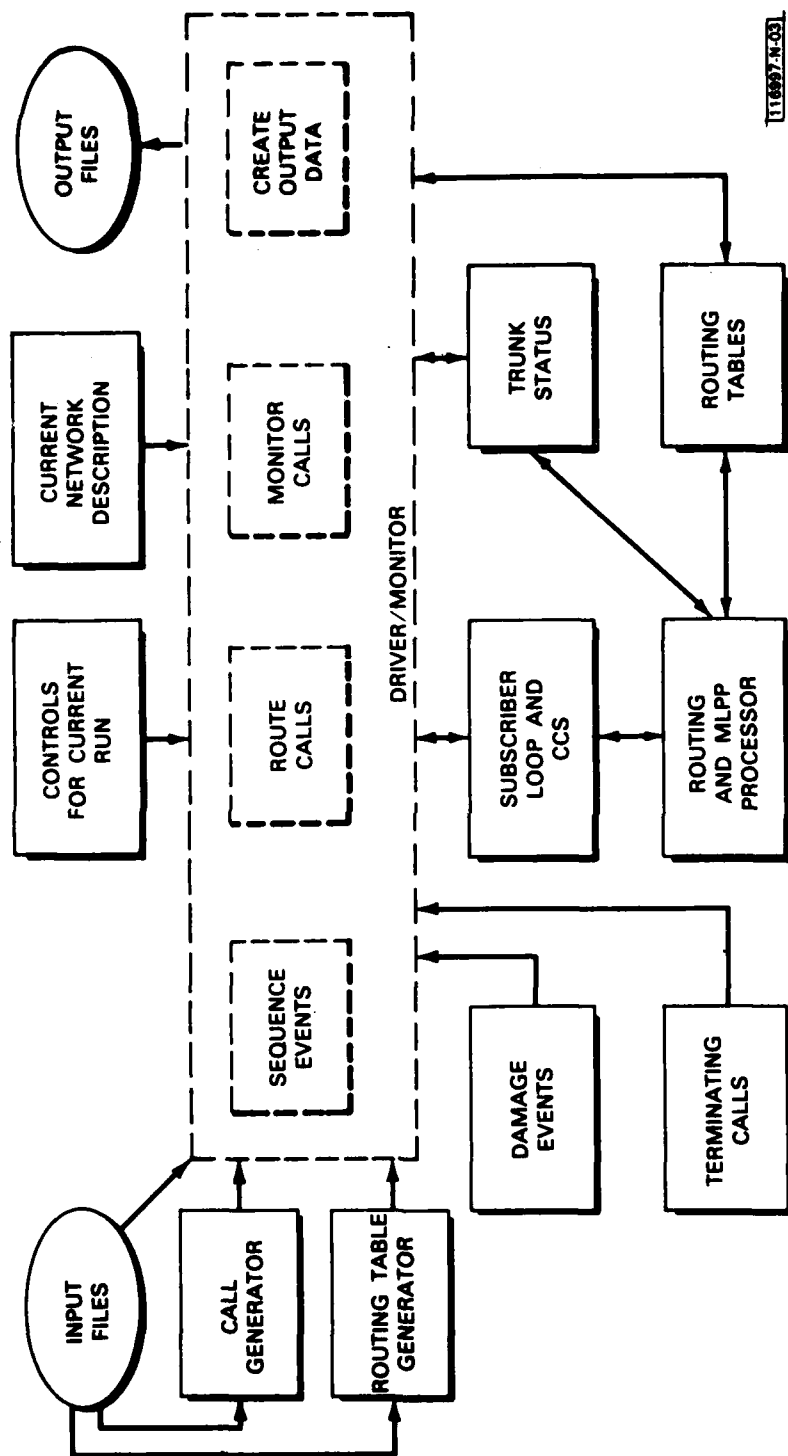
Fig. 11. Block diagram of call-by-call simulator.

30

statistics; (h) link CCS statistics; (i) point-to-point blocking probability distribution statistics and plots; and (j) time plots of total blocking, blocking by priority, number of calls preempted, call path length in links, CCS bits transmitted, number of low-priority calls terminated by each preempting high-priority call, percentage carried high-priority calls preempting to complete a call path, number of calls cranked back, and number of calls which queried remote earth stations. Plots are produced automatically from output files using a graphics software package called TELEGRAF.

Two sample output plots are presented in Figs. 12(a) and (b). Figure 12(a) is the average blocking for each precedence level vs time, and Fig. 12(b) is the number of calls preempted vs time. These plots include results for 1 h of simulated time using mixed-media routing with spill-forward control in network DSN1 when the satellite is destroyed at 30 min. Before the satellite is destroyed, blocking is below 0.02 and the number of calls preempted during 2-min. intervals is below 30. After the satellite is destroyed, blocking increases to roughly 0.2 for routine calls but stays at 0 for higher-precedence calls. The number of calls preempted increases after damage to roughly 110 calls every 2 min. because many high-precedence calls now must preempt to complete.

The simulator is currently being validated and thoroughly documented. We plan to add precedence flooding, source-destination preemption, and guided preemption during the first half of the next fiscal year. The simulator will then be used to compare all routing procedures, evaluate the different types of preemption, and also validate code in a routing/ control processor (RCP) to be interfaced to a real switch (see Sec. 4.4).
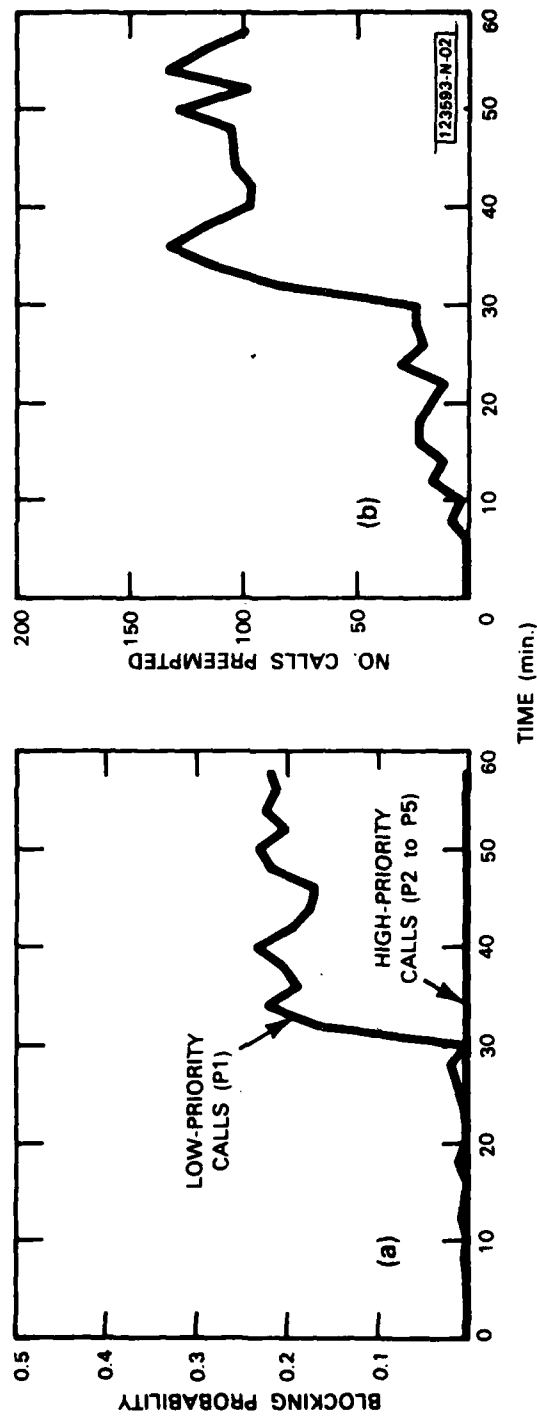
Fig. 12.   Sample call-by-call simulator output when the satellite in 20-node network DSN1 is destroyed at time = 30 min.   (a) Blocking probability by priority; (b) number of calls preempted.

32

3. EISN INSTRUMENTATION AND INTEGRATION

3.1 INTRODUCTION

The DSN is expected to include satellite connectivity overlaid on a network of digital circuit switches located on or near military bases. This satellite overlay will be designed to provide reduced cost for long-distance communications in normal conditions, as well as a degree of flexibility and endurability under limited stress conditions. A primary EISN program objective is demonstration of the satellite overlay concept in an experimental system with a limited number of network nodes. A major Lincoln effort has been the development of satellite/terrestrial interfaces and switch emulators to support these experiments. Detailed design and hardware/software implementation of the initial prototypes of this equipment was carried out in FY 81. The focus in FY 82 has been on the test and debugging, replication, deployment, and experimental application of these interfaces and emulators.

The EISN configuration, including photographs of the Lincoln-developed site equipment, is shown in Fig. 13. The system derives its satellite connectivity through the wideband satellite network (WB SATNET), a demand-assigned packet satellite network which supports both EISN experiments and DARPA-sponsored experiments in multi-user packet speech experiments. EISN equipment has been installed first at DCEC and Lincoln. The sites at the U.S. Air Force Rome Air Development Center (RADC) and at the U.S. Army Communications Command facilities at Ft. Huachuca, Arizona and Ft. Monmouth, New Jersey will become fully equipped during FY 83.

Experimental subsystems developed by Lincoln include: (a) a Packet/Circuit Interface (PCI) allowing access to the satellite channel from a circuit switch on a T1 carrier format digital interface, (b) a Telephone Office Emulator (TOE) which performs digital switch functions necessary for testing with the PCI, and (c) an Internet Packet Gateway (IPG) allowing connection between packet data networks and the WB SATNET for data protocol experiments. As indicated in Fig. 13, the PCI and IPG share common gateway hardware (a PDP-11/44 computer) for obtaining access to the WB SATNET. In addition, preliminary satellite/terrestrial routing experiments are controlled by means of software in the gateway. A major development during FY 82 has been the provision of terrestrial alternate routing (TAR) between PCIs by means of channel units attached to the TOE.

33

GATEWAY AND PACKET/
CIRCUIT INTERFACE

- PROTOCOL
  TRANSLATION

- DEMAND ACCESS TO
  SATELLITE TRUNKS

- SATELLITE/TERRESTRIAL
  ROUTING CONTROL

TO SATELLITE

EARTH
STATION
INTERFACE

PACKET
SATELLITE
DEMAND–
ASSIGNMENT
PROCESSOR

TO GATEWAY
AND TERRESTRIAL
NET

TO SATELLITE

TO PACKET
DATA NETWORKS

WIDEBAND
SATELLITE
NET

RADC

FT. MONMOUTH

DCEC

WESTAR III
(91 W)

TERRESTRIAL
TRUNKS

FT. HUACHUCA

TO TERRESTRIAL
TRUNKS

TELEPHONE
OFFICE EMULATOR

- SMALL CENTRAL
  OFFICE

- T1 INTERFACE TO
  SATELLITE TRUNKS

- CHANNEL UNITS
  TO TERRESTRIAL
  TRUNKS

Fig. 13.   Experimental integrated switched network (EISN) configuration.

34

Major FY 82 accomplishments in EISN instrumentation and integration have included:

(a) Integration, deployment, and test of PCIs and TOEs at DCEC and Lincoln;

(b) Construction of PCI and TOE hardware for the three additional EISN sites;

(c) Development and test of a TAR capability;

(d) Integration and test of IPGs with the Exploratory Data Network (EDN) at DCEC and with the ARPANET at Lincoln.

The status of equipment for the various EISN sites, and some FY 82 developments, are described in Sec. 3.2. Developments and improvements in the TOE are detailed in Sec. 3.3. The TAR capability is explained in Sec. 3.4. Finally, development of the IPG for packet data experiments on the EISN network is discussed in Sec. 3.5.

Section 3 focuses primarily on equipment development, with discussion of specific EISN experiments deferred to Sec. 4. The next phase of EISN instrumentation and integration will be to connect commercial circuit switches through T1 carriers to the PCIs, and to develop outboard routing/control processors for experiments in mixed-media routing. Section 4.4 describes design efforts with respect to this advanced routing/control experimental facility.

3.2 PACKET/CIRCUIT INTERFACE STATUS

Construction and development of additional TOE and PCI hardware for the sites other than Lincoln continued throughout FY 82. Highlights included the deployment of the PCI/TOE hardware to DCEC in February, and the addition of terrestrial routing in August. The initial operating capability for calls between TOE phones at Lincoln and DCEC, through PCIs and over the WB SATNET, was also demonstrated in February 1982. By the end of FY 82, all the boards for all five PCIs and TOEs had been built. A complete TOE chassis for the third site had been constructed and tested. The design for a multi-trunk echo canceler shelf that will be installed within the TOE chassis was completed and construction begun.

As experience was gained with the PCI, several minor refinements were introduced and retrofitted to units already built. It was found that the original power-up circuits did not always produce the correct initial state in the PCI processor boards. The use of a commercial power-up delay IC cured this problem. Two steps were taken to increase the throughput limit of the PCI. First, hardware improvements on the PCI processor board decreased the

delay involved in reading a byte from the channel bank. Second, a software update of the UMC-Z80 code in the PCI allowed local connections between T1 trunks (e.g., a call between two TOE phones) to be effected by a DMA controller on the UMC-Z80 rather than by a (more processor-intensive) packet route through the internet gateway. Such a direct connection is used when a call is routed over a dialed-up terrestrial trunk (Sec. 3.4). Occasionally, a momentary processing overload in the UMC-Z80 processor caused either a program crash or a clogging of the system with packets. This problem was relieved by a change in the real-time architecture of the UMC-Z80 code to shed processing load gracefully when necessary.

3.3 TELEPHONE OFFICE EMULATOR

The Telephone Office Emulator (TOE) was designed to look and act like a small class 5 telephone office. It has the capability of originating and receiving calls over a standard 1.544-Mbps T1 interoffice trunk line. The TOE was designed to provide an interim source of telephone traffic for testing the PCI equipment before connecting the PCI to an actual telephone switch. The prototype version was designed and constructed in FY 81. During FY 82, the debugging of hardware and software was completed. Construction of four more units has begun, and all boards are complete. The current version of the TOE will handle up to six subscriber loops. These loops use a modified phone which provides a 4-wire audio connection and separate signaling channels. The TOE can connect any of these phones to a selected channel on the T1 trunk line. Since the TOE uses special phones which provide 4-wire audio connections through to the T1 carrier, there is no need for echo suppression or cancellation.

The TOE currently supports either 7- or 10-digit dialing sequences, as does a standard telephone. If the second digit is a 0 or 1, then 10 digits are expected. Otherwise, 7 digits are expected. The TOE accumulates these digits internally. Upon completion of the dialing, the TOE seizes a trunk channel on the T1 carrier and pulses the number to the PCI using standard E and M signaling protocols. The TOE can also accept calls from the PCI. The PCI can seize a trunk into the TOE and then pulse the number using standard pulse dialing protocols. Only 7 digits are sent in this direction. The TOE examines the received number. If that number matches the number assigned to one of the local phones, then the TOE attempts to make the connection. If the phone is not busy, then a ringing signal is sent to the phone and the audio portion of the trunk is connected to a tone generator which provides the ringback to the calling party. When the called phone goes off hook, then the audio connection is completed to the phone. If the phone is busy, then the trunk is connected to a different tone generator which provides a busy signal to the caller.

36

In order to allow some terrestrial alternate routing experiments, the wiring of the TOE backplane was modified to allow two channel units to be connected directly to dedicated channels of the T1 carrier. These channel units are controlled directly from the PCI and will be used in various experiments involving direct connection to the switched telephone network.

Three TOEs have been completed and integrated with PCIs. One of these units was delivered to DCEC in February 1982; the others remain at Lincoln. All are currently in use in the EISN experimental program. The final two units are still under construction. All major pieces have been built. Integration and testing are under way.

3.4 TERRESTRIAL ALTERNATE ROUTING

One of the requirements of the DSN is the ability to operate after sustaining damage. Some of the routing experiments in the EISN test bed are designed to evaluate algorithms for routing around damaged portions of the network. In such experiments, the router has the choice of terrestrial trunks as well as the WB SATNET. While the full capability to conduct these experiments will come only with the introduction of commercial telephone switches and routing/control processors (Sec. 4.4) at the EISN sites, an interim capability for TAR was incorporated into the PCI during FY 82. The development of the interim TAR capability has provided some important benefits in addition to allowing initial routing experiments. These include overcoming the problems of interfacing EISN equipment to the public phone system, and introduction of the use of Common-Channel Signaling (CCS), in the form of datagrams over the WB SATNET, to control connection over a terrestrial trunk.

The key to obtaining a terrestrial trunk is giving the PCI the ability to dial calls through the PBX at its site. During FY 82, this was done by diverting two trunks of the T1 in the TOE channel bank from the control of the TOE processor and allowing their channel units to be used as ordinary channel units by the PCI, as shown in Fig. 14. A terrestrial trunk is dialed up for the duration of a routing experiment by placing a call through one of these channel units at one site to a similar channel unit at the other site via the public phone network. When the router decides to use this trunk to complete a call over the terrestrial trunk, it takes two actions:

(a) It loops the speech from the caller's T1 trunk to the T1 trunk leading to the terrestrial trunk.

CCS TO DISTANT PCI VIA SATELLITE

PACKET/
CIRCUIT
INTERFACE

T1

CHANNEL
UNITS
CONTROLLED
BY TOE
PROCESSOR

COMMON XMT/RCV EQPT

ORDINARY
CHANNEL
UNITS

TELEPHONE
OFFICE
EMULATOR

8-WIRE
TOE PHONES

PBX

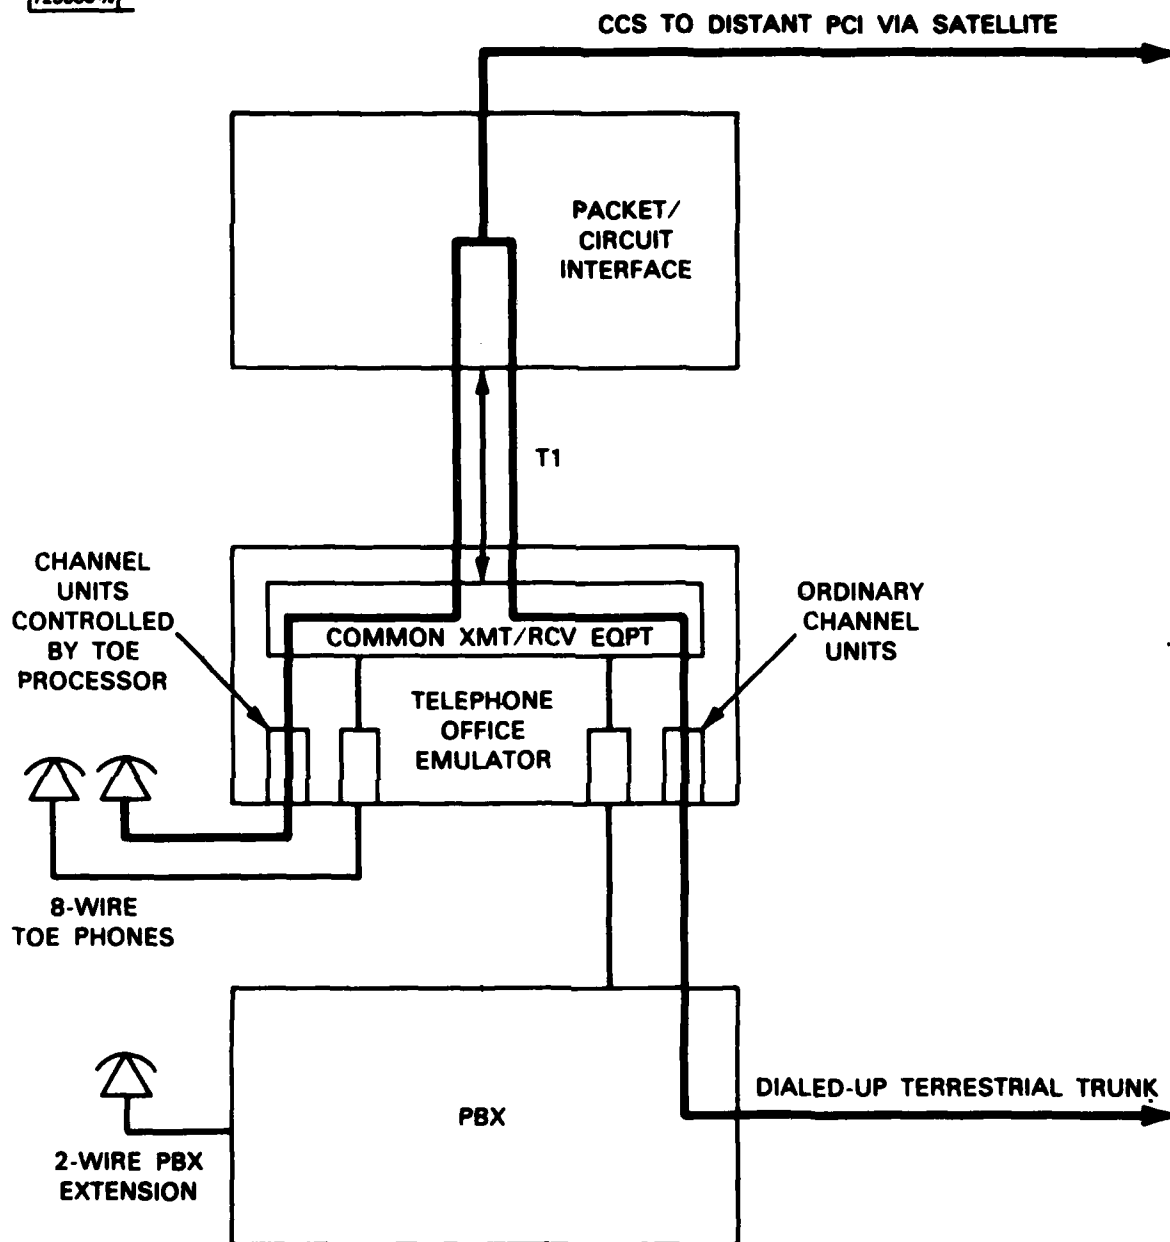DIALED-UP TERRESTRIAL TRUNK

2-WIRE PBX
EXTENSION

Fig. 14.   An EISN terrestrial trunk derived through TOE channel bank and PBX.

38

(*b*) It sends across the satellite channel a CCS datagram giving the phone number of the destination and the identifier of the terrestrial trunk to which the speech has been looped.

When one party to the terrestrially routed call hangs up, a CCS datagram from his end informs the PCI at the other end that the dialed-up trunk is now available for routing new calls.

The TAR capability has been demonstrated successfully between the Lincoln and DCEC sites. The PCI at each site contained a simple routing table, whose first choice was a satellite route and whose second choice was a TAR trunk. The PCI also stored the available capacity of the satellite route. For the purposes of this demonstration, that capacity was set to one call by a message typed into the internet gateway console. (After the automatic stream management software in the PSAT and the internet gateway comes into use, the routing algorithm can know the actual available satellite capacity.) Two calls were placed. The first was routed by satellite, while the second (finding the one satellite path already in use) was routed by the TAR trunk.

A side benefit of the terrestrial trunking faculty is the ability to call through the PCI to any phone on the PBX at the destination site, and even to local phones in the the public phone system around that site. During FY 83, 4-wire E&M tie trunks will be installed between the PBX and the PCI/TOE at most EISN sites. At that time it will become possible for any phone on the PBX at such a site to originate a call through the EISN network, not merely to receive one. Later, when the Stromberg-Carlson DBX-1200 circuit switch (see Sec. 4.4) is installed at Lincoln, similar 4-wire E&M tie trunks will connect it to the TOE/PCI to give it satellite connectivity. If the routing/control processor attached to the DBX-1200 should need the full capacity of the PCI, the TOE could be disconnected [as in Fig. 15(a)] and an ordinary digital channel bank used to connect the trunk interface of the DBX-1200 to the PCI. Later, when a direct T1 interface for the DBX-1200 becomes available from the manufacturer, a .... digital connection [Fig. 15(b)] can be used.

Now that the PCI has been connected to the public phone system, there are possibilities for echo from the 2- to 4-wire transitions. Because the delays caused by packet processing can be over 100 ms, even physically short paths can suffer perceptually disturbing echo. Therefore, echo cancelers from COMSAT General Telesystems have been installed in all voice paths between the PCI and the PBX at its site. The operation of the echo canceler is
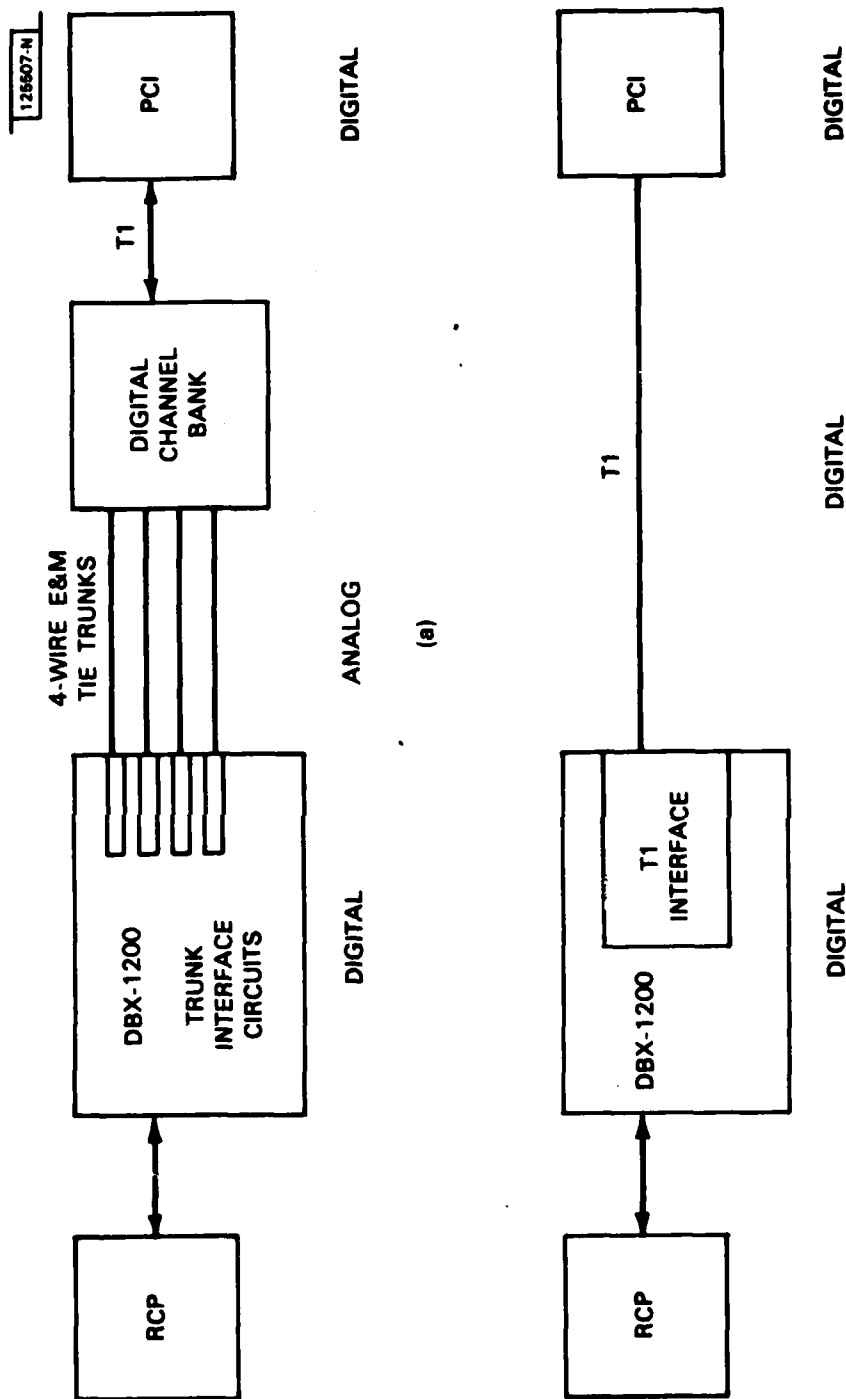
39

Fig. 15. PCI interfaced to switch through (a) analog trunks and ordinary channel bank, and (b) all-digital T1 interface on switch.
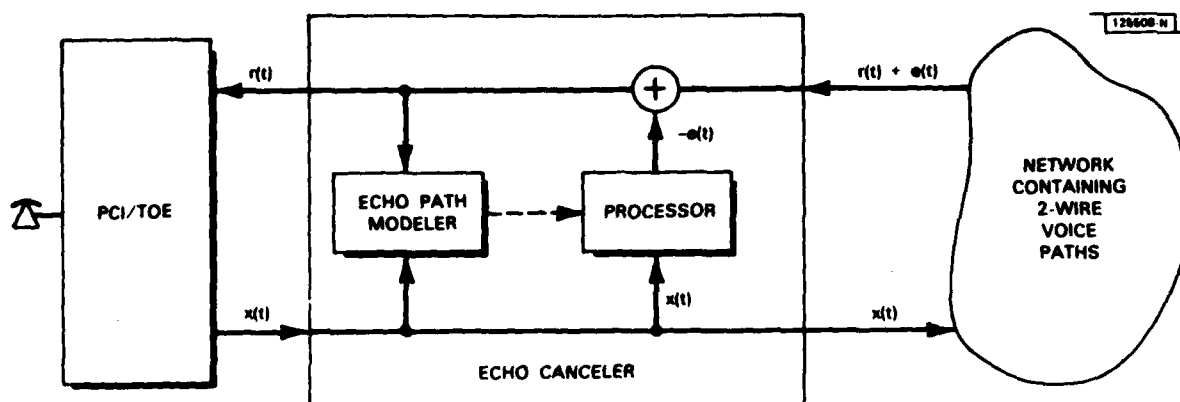
40

Fig. 16.   Operation of echo canceler in EISN experiments.

shown in Fig. 16. A speech waveform x(t), transmitted from the PCI into a phone network
with 2-wire paths, may be reflected and return as an echo e(t), added to the speech r(t)
received from the phone network. During the silences in r(t), the echo canceler tries to ana-
lyze, and duplicate within its speech processor, the echo path that transforms x(t) into e(t)..
The estimated echo is then subtracted from r(t) + e(t), leaving (in the ideal case) only r(t) to
enter the PCI. Although analysis occurs only during silences in r(t), echo cancellation occurs
continuously. Therefore, the echo canceler, unlike an echo suppressor, functions properly
even when one speaker is interrupting the other. The use of the echo cancelers has rendered
initially severe echo practically imperceptible.

3.5   INTERNET PACKET GATEWAY

To support the data protocol performance measurement experiments to take place in
FY 83, software for the PDP-11/44 gateway at DCEC has been extended to provide a con-
nection to EDN and to handle the Source Route options specified in the DoD standard
Internet Protocol (IP). The gateway hardware consists of a PDP-11/44 computer with three
UMC-Z80 interface processors. The software is modular, and modules are available to sup-
port connections to the WB SATNET, the PCI, a Lincoln-developed local-area packet voice
network (LEXNET), and EDN. The EDN software is the same as that used at Lincoln to
provide a gateway to ARPANET. The DCEC gateway has three interface processors and
can be configured to interconnect any three of the above-mentioned networks at any one
time. The UMC-Z80s associated with the WB SATNET, PCI, and EDN are internally con-
nected to additional interface equipment associated with those networks. Connection of the

41

LEXNET involves an external cable which currently can be plugged to either the UMC-Z80 that serves the PCI or the one that serves the EDN, but not to the one that serves the WB SATNET. When the LEXNET is in use, the net normally associated with the chosen UMC-Z80 cannot be served by the gateway. As a result, the DCEC gateway can have the following three configurations:

(a) WB SATNET, PCI, EDN

(b) WB SATNET, PCI, LEXNET

(c) WB SATNET, LEXNET, EDN.

Configurations (a) and (c) are of interest for data experiments.

The 11/44 gateway supports both the datagram internet protocol (IP) and the stream protocol (ST) used for packet speech transmission. The higher-level network voice protocol (NVP-II) makes use of IP datagrams for point-to-point and conferencing setup messages, but the voice itself is handled by ST. To support data protocol experiments, we have extended the IP/ST gateway capabilities to include features of IP not required for voice use. Among these features are the IP Source Route and Route-Record options. The Source Route option allows the sender of a datagram to specify an arbitrary route through a slice of IP gateways. The Route-Record provides a path history that can be used by the receiver of a datagram to provide a reverse path for a reply or acknowledgment. Source-routing is an important capability needed for protocol performance experiments, since it allows a wide variety of paths with different delay and throughput characteristics to be explored from a small number of origin and destination experiments sites. Without it, only the shortest path would be used by the packets; and it would be necessary to move to a variety of locations to explore other paths.

Software to support the Source Route and Route-Record options is now operational in the IP/ST gateway. It was used to test the operation of the EDN connection by looping source-routed IP datagrams from Lincoln over the WB SATNET, through the IP/ST gateway, across EDN to the IP gateway between EDN and ARPANET, and back over the same path.

4. EISN EXPERIMENT PLANNING, COORDINATION, AND EXECUTION

4.1 EXPERIMENT PLANNING

Over the next several years, the hardware and software configurations at the EISN nodes will be used as a scaled experimental test bed for addressing and resolving system design issues for the DSN. The design and implementation of the test bed itself is an ongoing process which is highly interactive with the plans and objectives for its use. Certain capabilities are now in place and others are programmed for the near future, and in every case the design of successive augmentations and changes must be preceded by carefully thought out experiment plans which motivate them. Similarly, the spectrum of experiments achievable at a given time is determined by the capabilities currently available. Thus, a major part of Lincoln's work in this program is the development and continual evolution of experiment plans directed toward DSN objectives.

In November 1981, the EISN Experiment Plan was published as Lincoln Laboratory Project Report NSST-2. It began with a review of the development of the DSN concept in the DoD, including a statement of the desired attributes of the DSN. The relationship of planned EISN experiments to DSN system issues was discussed in detail. An overview of the EISN test-bed configurations was given, and a phased implementation schedule extending over several years was set up. The document then gave detailed descriptions of the purpose, background, evaluation criteria, and test plans for three general categories of scaled DSN experiments — namely routing, flow control, and system control; satellite and terrestrial integration; and system security. The EISN Experiment Plan was followed by a Work Plan detailing the specific items that were to be carried out during FY 82. The status of these items is discussed below in Sec. 4.3. A similar Work Plan is currently in preparation for FY 83, reflecting many of the advanced experimental concepts described below.

The EISN Experiment Plan detailed a sequence of specific experiment categories organized in terms of chronological development of the test bed, beginning with a basic equipment configuration installed at two sites. This configuration consists of: the satellite earth station; the high-speed modem and coder-decoder, called the Earth Station Interface (or ESI); the channel demand-assignment processor, called the Pluribus Satellite Interface Message Processor (or PSAT); the Lincoln-built Packet/Circuit Interface (or PCI), which has a T1 interface compatible with conventional telephone switching equipment; and the

Fig. 17. Planned EISN site equipment configurations.

Lincoln-built Telephone Office Emulator (or TOE), which provides an interim mechanism for combining the signals from a few telephones into a T1 carrier as well as a limited switching capability. Such a configuration was available at two sites (DCEC and Lincoln) during FY 82, and a corresponding set of two-site experiments was laid out in the Plan. The second phase of test-bed development, to take place during FY 83, was installation of the basic equipment configuration at the three MILDEP EISN sites (Fts. Monmouth and Huachuca, and Rome Air Development Center). A corresponding set of multisite experiments was detailed in the Plan.

The third phase of test-bed development described by the EISN Experiment Plan was incorporation of an "off-the-shelf" digital circuit switch and an associated outboard processor into the test-bed equipment at each of the sites, together with an experimental common-channel signaling (CCS) network and mechanisms for linking EISN with a large-scale call-by-call network simulator. The planned configuration is illustrated in Fig. 17. It was indicated that each of the experiments done with the basic equipment would be repeated after the upgrade, with significant corresponding increases in realism and in potential for resolving important DSN issues.

Several developments have occurred this year which have given structure and impetus to both test-bed design and advanced experiment planning, with respect to both implementation and use of full site configurations. One is the rapid evolution that has occurred in planning and preparation for the DSN itself, in the DSN Program Office at the DCA. Another is the increased involvement of the MILDEP site organizations in the EISN experiment and its planning. A third factor has been crystallization of ideas and plans for the actual implementation of switch interconnections and RCPs at the various sites. A decision has been made to procure and install a medium-sized dedicated switch at both Lincoln and DCEC, together with PDP-11/44 RCPs, as described below under the heading of Advanced Routing/Control Experimental Facility Design (Sec. 4.4). The switch and processor for Lincoln are on order, with delivery expected early in FY 83; the DCEC switch and processor are to be ordered as soon as possible, and they will be initially delivered to Lincoln for integration and for preliminary two-switch experiments. Further details are given below.

4.2 EXPERIMENT COORDINATION

During this year, Lincoln has undertaken an expanded role in satellite system engineering for EISN, in the sense of coordinating the integration of facilities at the various sites as well as network-wide coordinating of experiments. One example of this activity is the plan

45
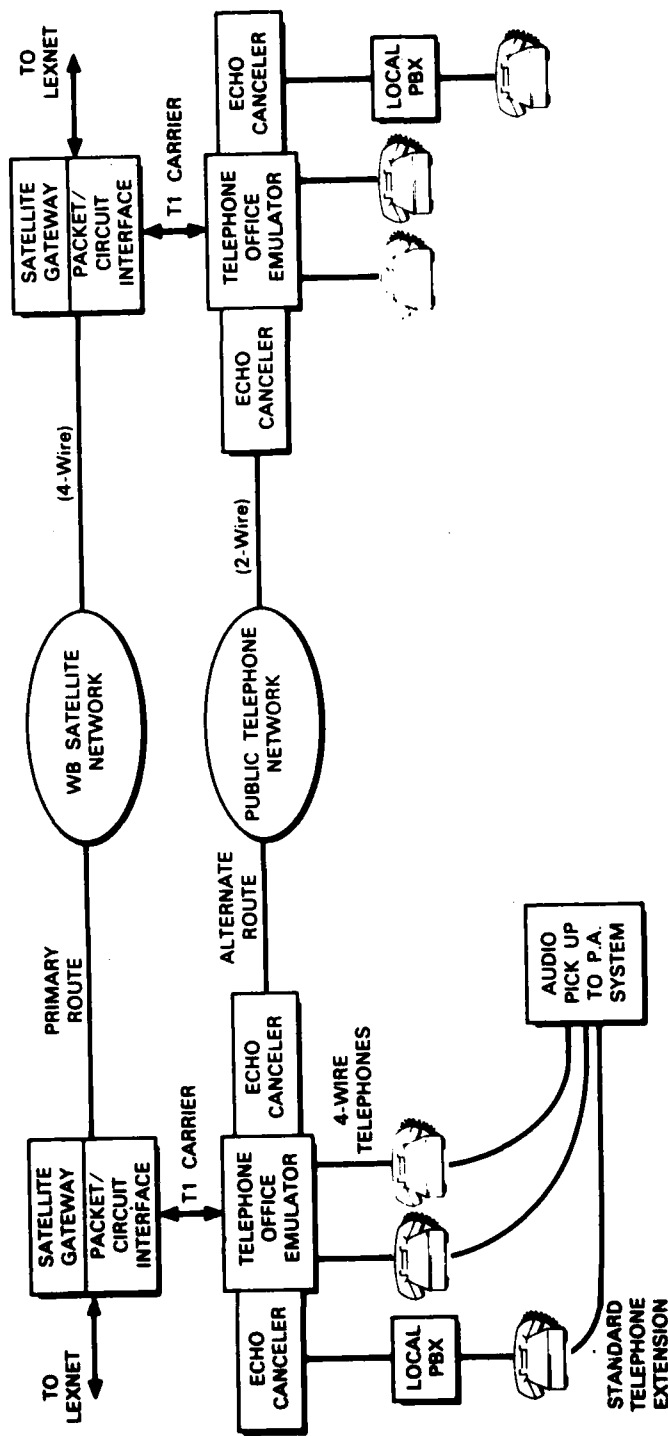
Fig. 18. Generic two-site EISN experiment configuration.

46

currently being implemented for Lincoln to become the frequency and power-level reference station for the network. This involves installation at Lincoln of a spectrum analyzer with an accurate internal frequency reference, and cooperative activity with BBN in devising ways to measure and compare transmitted signal parameters from each network site. This approach applies both to maintaining the calibration of existing stations and assisting in the installation of new ones. Another example of this expanded coordination activity is the role played by Lincoln in the June 1982 packet speech demonstration, as described below in Sec. 4.3, Experiment Milestones.

At the start of FY 82, both Lincoln and DCEC had their EISN satellite network equipment, but neither had an operational TOE or PCI. A major component of our work has been the construction and integration of this equipment, not only for Lincoln and DCEC but for the three MILDEP sites. After delivery of DCEC's equipment in the second quarter of the fiscal year, it became physically possible to carry out experiments as detailed in the Work Plan, and the status and completion of these experiments is described below. This experimental activity culminated in a program review and demonstration, held at DCEC on 7 October 1982, which essentially repeated the major experiments, as described below.

4.3 EXPERIMENT MILESTONES

During FY 82, the focus of the EISN experimental effort was on two-node experiments in satellite/terrestrial integration, routing, and internet data communication. These experiments focused on test and utilization of the PCI, TOE, and IPG equipment at the Lincoln and DCEC sites. Figure 18 shows the generic EISN configuration which has been used for experiments in multi-media call setup and satellite/terrestrial alternate routing, the planning for which is described in the FY 82 Work Plan.

Multi-media call setup represents the capability to dial a call over either a satellite or terrestrial routing. The initial cross-satellite calls between Lincoln and DCEC were carried out in February 1982. Three simultaneous full-duplex 64-kbps connections have been maintained. Some modification in the satellite DAMA frame allocation is needed to increase this to four, which represents the throughput we expect to be achievable with a single PCI. In March 1982, a call between TOE phones at Lincoln and DCEC was conducted simultaneously with a call between LEXNETs at Lincoln and ISI. This exercised the capability of the use of satellite overlay on a terrestrial circuit-switched net, where the demand-assigned satellite channel is shared with other users. The capability to set up calls via a terrestrial route was established at Lincoln in August, and tested between Lincoln and DCEC in September.

47

As described in Sec. 3.4, the satellite/terrestrial alternate routing experiments make use of a capability called Terrestrial Alternate Routing (TAR), under which a long-distance call between the two PCIs is dialed up and held for the duration of the experiments, during which it is treated as though it were a trunk, thus avoiding the excessive cost of leasing full-period lines. The alternate routing experiment makes use of an interim table-driven route selection module in the PCI which places the first of two calls over the satellite channel, and the second over a land line. The setup shown in Fig. 18 has also been used to support dialing area extension (DAX) experiments, where calls are made between TOE phones and standard telephone extensions on local PBXs at Lincoln and DCEC. In addition, preparations are being made for installing a basic preemption capability in the PCI software, and conducting experiments involving preemption of calls routed either terrestrially or via satellite.

Data internetting objectives, as set forth in the Work Plan, primarily consisted of test and validation of the EDN/EISN gateway (IPG). Initial experiments have been conducted using a Lincoln-built measurement host to collect delay and delay-dispersion results for a data internet environment including the EDN, the ARPANET, and the WB SATNET. The data internet configuration is shown in Fig. 19. Using source-routing, packets have been
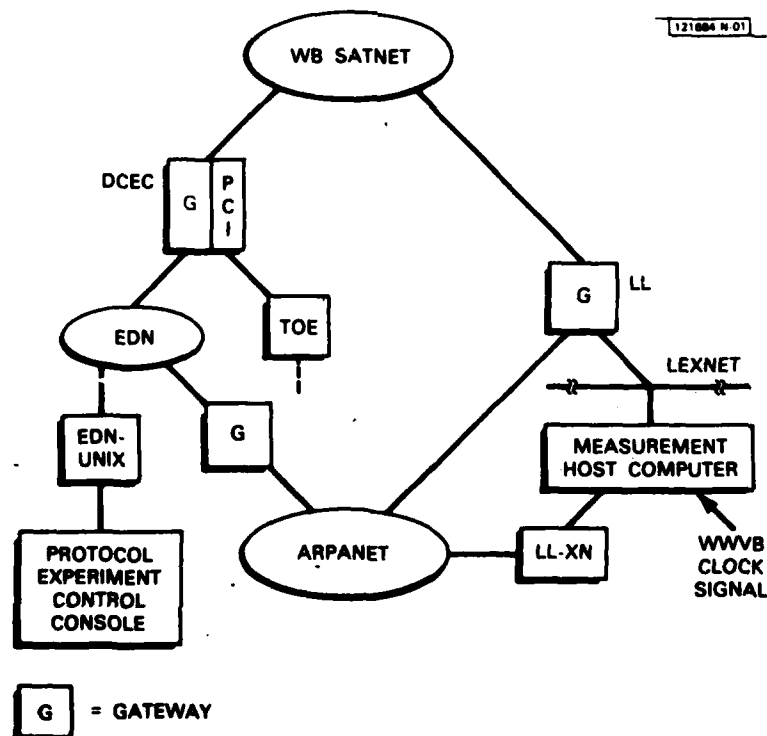

Fig. 19. DCEC/Lincoln data protocol experiment configuration.

48

directed over various loop paths involving the networks shown. Delay histograms are collected in the measurement host. Figure 20 is a typical histogram of the delay-dispersion results over a path consisting of a loop from the measurement host to the EDN and back via the ARPANET. There is no restriction to loop paths, as a standard global clock signal (WWVB) is available for precise delay measurements over one-way network paths.

A major demonstration of DARPA-oriented WB packet speech equipment was held at Lincoln Laboratory on 3 June 1982, with excellent success. While this demonstration did not directly exercise DCA facilities such as the Packet/Circuit Interface, the required coordination efforts were of great importance in integrating the WB SATNET equipment which serves as a communications medium for both DARPA and DCA experiments. The demonstration exploited packet speech local network and terminal equipment at the Information Sciences Institute of the University of Southern California (Los Angeles) and at SRI International (San Francisco Bay Area), as well as at Lincoln Laboratory. During preparation for the demonstration, a local-area packet voice network (LEXNET) was installed at DCEC, so that DCEC could serve as an alternate packet voice site. Later, the DCEC LEXNET was used for circuit/packet interoperability demonstrations, where calls were established across the WB SATNET between TOE phones at Lincoln or DCEC, and LEXNET voice terminals at the other site.

4.4 ADVANCED ROUTING/CONTROL EXPERIMENTAL FACILITY DESIGN

During FY 82, detailed design efforts have begun on the advanced routing/control experimental facility (Fig. 17) including commercial switches and routing/control processors (RCPs). Switch selection for the Lincoln and DCEC sites has been made, as described below. RCP implementation and integration of the RCP/switch facility are major planned efforts for FY 83.

The basic idea of the RCP design is that any of a variety of modern computer-controlled switches can be enhanced by the addition of an outboard processor to perform sophisticated routing, control, and MLPP functions, exploiting CCS connections to similar processors on other switches in the network. Figure 21 shows the concept.

In the EISN context, several factors have entered into the selection of the particular switch types that can be interfaced with RCPs at the various nodes. The choice is quite straightforward for Ft. Huachuca, where a new Stromberg-Carlson DBX-5000 has recently been installed to support a major share of the operational traffic on the base. Largely
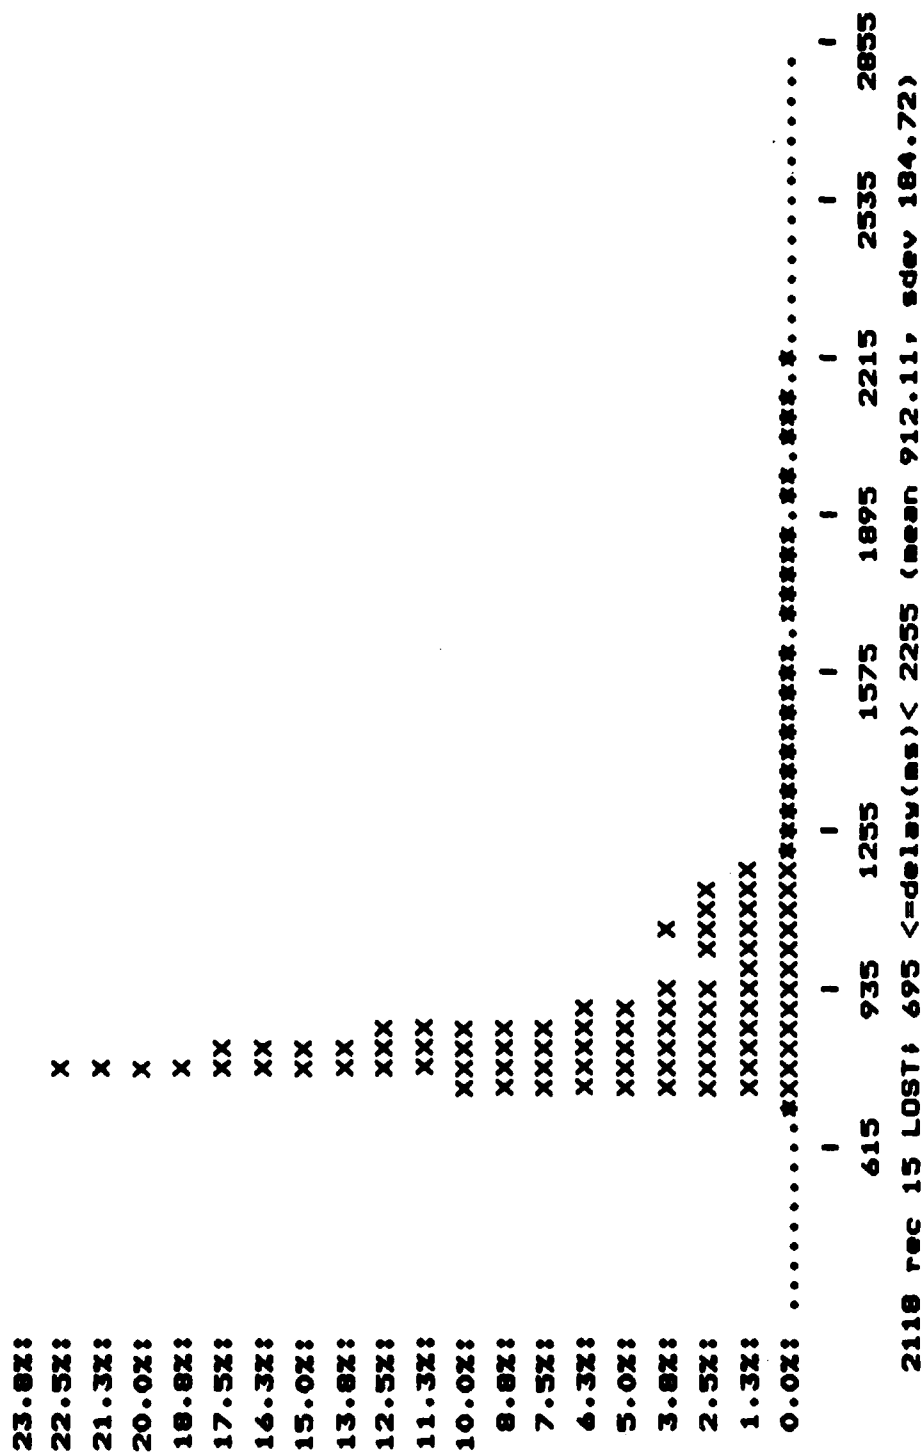
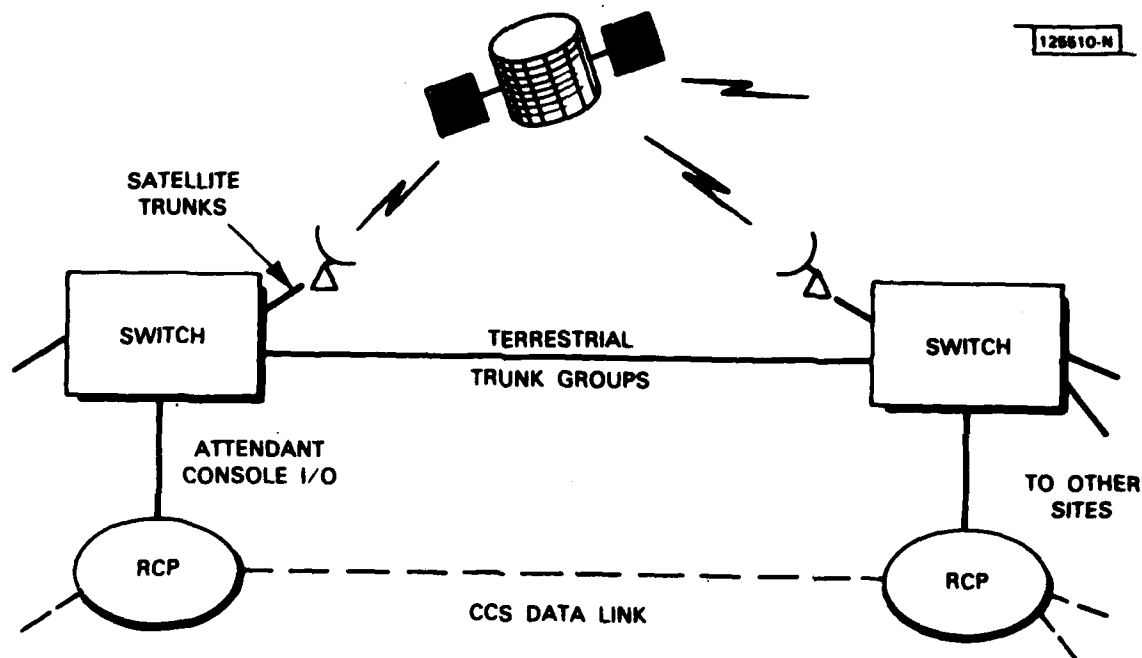Fig. 20. Typical packet delay histogram for ARPANET/EDN loop path.

Fig. 21.   RCP/switch experimental facility.

because of this factor, it was decided that Lincoln would purchase a dedicated experimental switch of identical architecture, but smaller in size by a factor of four (the Stromberg-Carlson DBX-1200), and would develop the RCP and interfaces under controlled experimental conditions before replicating and installing them at Ft. Huachuca. The Lincoln switch is on order, with delivery expected early in FY 83. A second DBX-1200 will be purchased and installed at DCEC as a dedicated experimental facility, with an RCP to be provided by Lincoln. Ft. Monmouth has a Northern Telecom SL-1 switch serving the base's operational needs, and it appears that it would be appropriate to design the necessary interface to allow an RCP to be integrated with it. Finally, RADC is expected to be obtaining a SCOPE DIAL switch, which will be the Northern Telecom DMS-100; this is a reasonable candidate for integration with an RCP.

Figure 22 illustrates the architecture of the Stromberg-Carlson DBX, which is typical of modern switches. The key to the RCP interface is the remote attendant console, which is connected directly to the switch CPU by voice and data lines, and has an array of standard control functions inherent in the off-the-shelf software structure of the switch. The RCP will
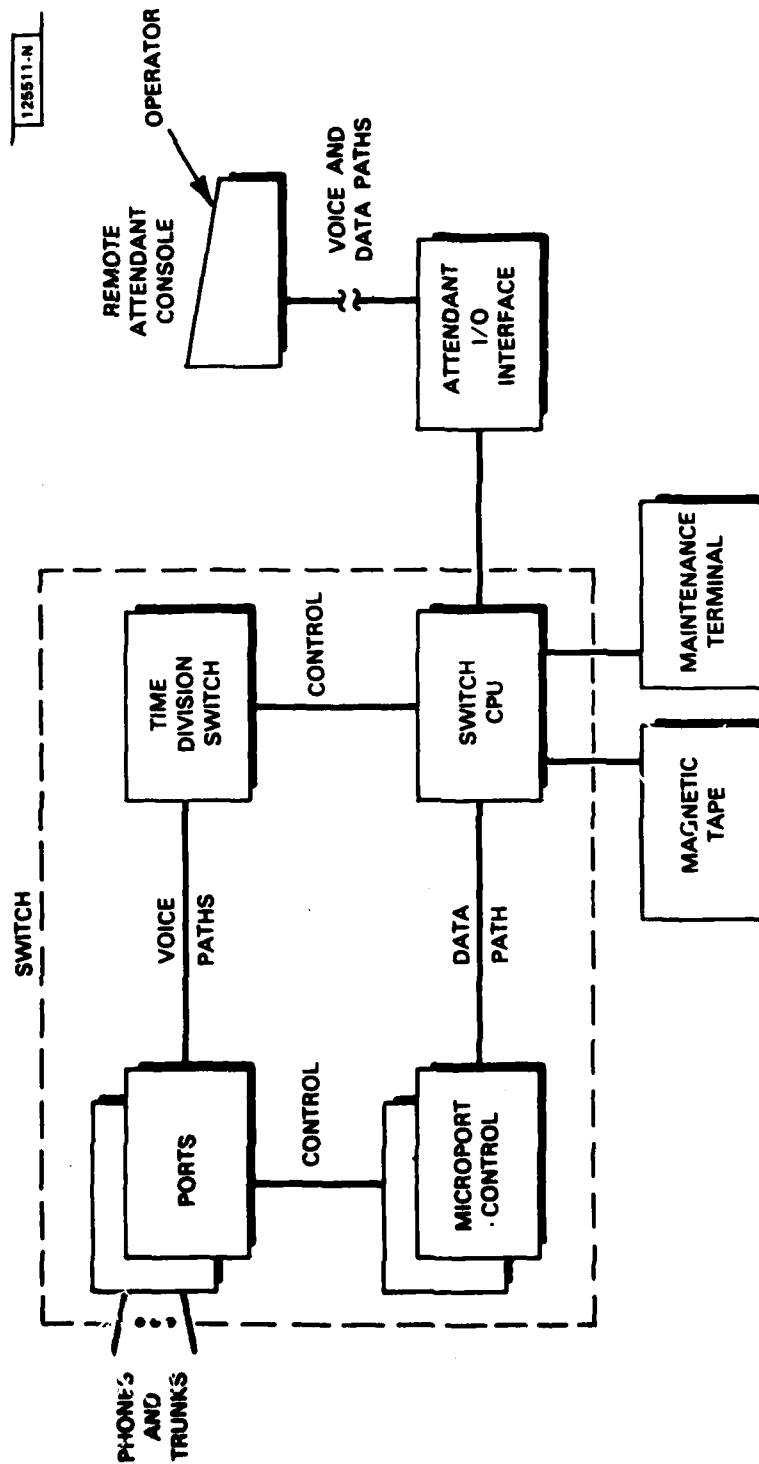
51

Fig. 22. Stromberg–Carlson DBX architecture.

52

replace the attendant console, using special Lincoln-designed interface modules which will convert attendant I/O signals to and from ASCII characters.

As noted above, the Lincoln switch will be delivered early in FY 83. RCP software design is already in progress, and the basic RCP/switch capabilities are expected to be put in place and checked out during FY 83. The plan is to have the DCEC switch delivered to Lincoln initially, for integration with its RCP and for convenient execution of a group of two-node experiments.

REFERENCES

1. Network Speech Systems Technology Program Annual Report, Lincoln Laboratory, M.I.T. (30 September 1981), DTIC AD-A112718/2.

2. R.P. Lippmann, "Steady State Performance of Survivable Routing Procedures for Circuit-Switched Mixed-Media Networks," Technical Report 633, Lincoln Laboratory, M.I.T. (in press).

3. W.C. Miller, T.D. Walsh, J. Van Zuylen, A. Covo, and T.G. Huemiller, "Multi-Level Precedence and Preemption Study," Defense Communications Engineering Center Engineering Publication No. A04-81 (November 1981).

4. J. McQuillan, I. Richer, and E.C. Rosen, "The New Routing Algorithm for the ARPANET," IEEE Trans. Commun. *COM-28*, 711-719 (1980).

5. G.W. Bernas and D.M. Grieco, "A Comparison of Routing Techniques for Tactical Circuit-Switching Networks," in Proc. 1978 Intl. Conf. on Communications (ICC '78), Toronto, Canada, 4-7 June 1978, pp. 23.5.1-23.5.5

6. C.N. Shearer, "Modified Forward Routing for CONUS/CSN AUTOVON," Defense Communications Engineering Center Technical Report 26-75 (September 1975).

7. M.J. Fischer, D.A. Garbin, T.C. Harris, and J.E. Knepley, "Large Scale Communication Networks — Design and Analysis," Omega *6*, 331-340 (1978).

8. S. Katz, "Statistical Performance Analysis of a Switched Communication Network," in Proc. 5th Intl. Teletraffic Congress, New York, June 1967.

9. M.J. Fischer and J.E. Knepley, "Circuit-Switched Network Performance Algorithm (MOD-1)," Defense Communications Agency Technical Note 1-73 (January 1973).

GLOSSARY

CCS   Common-Channel Signaling
CPU   Central Processing Unit

DAMA   Demand-Assignment Multiple Access
DCA   Defense Communications Agency
DCEC   Defense Communications Engineering Center
DMA   Direct Memory Access
DSN   Defense Switched Network

EDN   Exploratory Data Network
EISN   Experimental Integrated Switched Network
ESI   Earth Station Interface

IP   Internet Protocol
IPG   Internet Packet Gateway

MLPP   Multi-Level Precedence and Preemption

NVP   Network Voice Protocol

PCI   Packet/Circuit Interface
PSAT   Pluribus Satellite Interface Message Processor

RADC   Rome Air Development Center
RCP   Routing/Control Processor

SATNET   Satellite Network
ST   Stream Protocol

TAR   Terrestrial Alternate Routing
TOE   Telephone Office Emulator

WB SATNET Wideband Satellite Network

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER ESD-TR-82-167	2. GOVT ACCESSION NO. *A128 932*	3. RECIPIENT'S CATALOG NUMBER
4. TITLE *(and Subtitle)* Defense Switched Network Technology and Experiments Program		5. TYPE OF REPORT & PERIOD COVERED Annual Report 1 October 1981 - 30 September 1982
		6. PERFORMING ORG. REPORT NUMBER
7. AUTHOR(s) Clifford J. Weinstein		8. CONTRACT OR GRANT NUMBER(s) F19628-80-C-0002
9. PERFORMING ORGANIZATION NAME AND ADDRESS Lincoln Laboratory, M.I.T. P.O. Box 73 Lexington, MA 02173-0073		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS Program Element No. 33126K
11. CONTROLLING OFFICE NAME AND ADDRESS Defense Communications Agency 8th Street & So. Courthouse Road Arlington, VA 22204		12. REPORT DATE 30 September 1982
		13. NUMBER OF PAGES 68
14. MONITORING AGENCY NAME & ADDRESS *(if different from Controlling Office)* Electronic Systems Division Hanscom AFB, MA 01731		15. SECURITY CLASS. *(of this report)* Unclassified
		15a. DECLASSIFICATION DOWNGRADING SCHEDULE

16. DISTRIBUTION STATEMENT *(of this Report)*

Approved for public release; distribution unlimited.

17. DISTRIBUTION STATEMENT *(of the abstract entered in Block 20, if different from Report)*

18. SUPPLEMENTARY NOTES

None

19. KEY WORDS *(Continue on reverse side if necessary and identify by block number)*

routing algorithms	Defense Switched Network	circuit-switched network
Experimental Integrated	mixed-media routing	packet switching
Switched Network	pre-emption	

20. ABSTRACT *(Continue on reverse side if necessary and identify by block number)*

This report documents work performed during FY 1982 on the DCA-sponsored Defense Switched Network Technology and Experiments Program. The areas of work reported are: (1) development of routing algorithms for application in the Defense Switched Network (DSN); (2) instrumentation and integration of the Experimental Integrated Switched Network (EISN) test facility; and (3) EISN experiment planning, coordination, and execution.